

Architecture and Concept of Operations for a Warfighter's Internet

Volume 1



Prepared for:
Defense Advanced Research Projects Agency
Information Systems Office

JPL

MITRE

Edited by:
Massachusetts Institute of Technology
Lincoln Laboratory

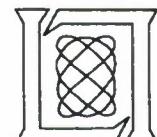


28 January 1998



Prepared for the Defense Advanced Research Projects Agency under
Air Force Contract F19628-95-C-0002

Approved for public release; distribution is unlimited.



ADA336795

This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. The work was sponsored by the Defense Advanced Research Projects Agency, ISO, under Air Force contract F19628-95-C-0002.

This report may be reproduced to satisfy needs of U.S. Government agencies.

The ESC Public Affairs Office has reviewed this report, and it is releasable to the National Technical Information Service, where it will be available to the general public, including foreign nationals.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER



Gary Tutungian
Administrative Contracting Officer
Contracted Support Management

Non-Lincoln Recipients

PLEASE DO NOT RETURN

Permission is given to destroy this document
when it is no longer needed.

Architecture and Concept of Operations for a Warfighter's Internet

Volume 1

Prepared with Contributions from:

**Air Force Rome Laboratory
Army CECOM
Jet Propulsion Laboratory
MIT Lincoln Laboratory
MITRE Corporation
Navy Command Control Ocean Surveillance Center
Naval Research Laboratory**

Prepared for:

**Defense Advanced Research Projects Agency,
Information Systems Office**

28 JANUARY 1998

Approved for public release; distribution is unlimited.

TABLE OF CONTENTS

List of Illustrations	v
List of Tables	vii
1. OVERVIEW AND EXECUTIVE SUMMARY	1-1
1.1 Introduction	1-1
1.2 Operational Scenarios and Communications Needs Overview	1-2
1.3 Architectural Concept and Design Overview	1-6
1.4 Summary of Characteristics of the Warfighter's Internet	1-10
2. OPERATIONAL NEED AND CONCEPT	2-1
2.1 Operational Scenarios and Communications Needs	2-1
2.1.1 Overview	2-1
2.1.2 Army Tactical Communications Needs	2-1
2.1.3 Navy Tactical Communications Needs	2-5
2.1.3.1 Today's Needs	2-5
2.1.3.2 Evolving Needs	2-7
2.1.3.3 Goal Properties	2-11
2.1.3.4 Warfighter's Internet Implications	2-13
2.1.4 Air Force Tactical Communications Needs	2-14
2.1.5 Warfighter's Internet Applications in T&E and Training	2-15
2.1.5.1 Introduction	2-15
2.1.5.2 Test & Evaluation (T&E)	2-15
2.1.5.3 Training	2-19
2.1.5.4 Summary of Potential T&E Training Applications of the WI	2-22
2.2 Information Architecture	2-22
3. ARCHITECTURAL CONCEPT AND DESIGN	3-1
3.1 Introduction / Overview	3-1
3.2 Interfaces, Boundary of Warfighter's Internet	3-2
3.3 Efficient Communications	3-4
3.4 Communication Element Design Considerations	3-6
3.4.1 Uplinks and Downlinks to Individual Subscribers	3-7
3.4.1.1 Introduction	3-7
3.4.1.2 Design Criteria	3-8
3.4.1.3 Uplink and Downlink Performance and Discussion	3-11
3.4.2 Uplinks and Downlinks to Terrestrial Nodes	3-16
3.4.3 Airborne Node to Airborne Node Cross-links	3-17
3.4.3.1 Directive Beams	3-18
3.4.3.2 Omnidirectional Beams	3-20
3.4.3.3 Topology Considerations	3-21

TABLE OF CONTENTS (Continued)

3.5	Functional Architecture for the Mobile WI	3-23
3.5.1	General Architecture and Design Goals	3-23
3.5.2	Warfighter's Internet Architecture (Commonalities and Variations)	3-26
3.5.2.1	Overview	3-27
3.5.2.2	Summary of Architecture Option 1	3-37
3.5.2.3	Summary of Architecture Option 2	3-44
3.5.3	Entry Node	3-55
3.6	Addressing and Routing	3-56
3.6.1	IP Address Structures	3-56
3.6.2	IP Addresses and Mobility	3-57
3.6.3	IP Addresses and Multicasting	3-60
3.7	Security	3-61
3.7.1	Security and IP (Historic)	3-63
3.7.2	Voice Security (Historic)	3-64
3.7.3	Key Management	3-64
3.7.4	IPv6 and Security	3-65
3.7.5	The Potential Role of Fortezza	3-69
3.8	Network and Systems Management	3-71
3.8.1	Management Protocols and Architecture	3-71
3.8.1.1	Implementation Schemes	3-72
3.8.2	Network Management Functions	3-73
3.8.2.1	Fault Management	3-73
3.8.2.2	Configuration Management	3-73
3.8.2.3	Performance Management	3-73
3.8.2.4	Security Management	3-74

APPENDICES (VOL. 2)

Appendix A: Cross-Links Segment Description (Architecture Option 2)	A-1
Appendix B: Protocols to Support Integrated Services in a Mobile Cellular System Architecture (for Architecture Option 2)	B-1
Appendix C: Shared Resources Multiple Access (SRMA)	C-1
Appendix D: ATM Integration for Architecture 2	D-1
Appendix E: WI Architecture Description (Option 1)	E-1
Appendix F: Multicast Concepts for Warfighter's Internet	F-1
Appendix G: Frequency Management and Coordination	G-1
Appendix H: Assessment of Candidate Protocols for Warfighter's Internet	H-1
Appendix I: Derivation of Optimum Packet Length for ARQ Wireless Data Communications	I-1
Appendix J: Uplink and Downlink Frequency Assignments and Reuse	J-1
Appendix K: Warfighter's Internet Functional Specification for Army Tactical Communications	K-1
Appendix L: Topology and Electromagnetics Systems Management	L-1

LIST OF ILLUSTRATIONS

Figure No.		Page
1-1	Warfighter's Internet concept.	1-2
1-2	Warfighter's Internet provides support of military operations.	1-3
1-3	Warfighter's Internet as overlay to existing tactical communications.	1-4
1-4	Integrated information flow in the battlefield.	1-5
1-5	Efficient use of uplink and downlink resources.	1-7
1-6	Elements of the Warfighter's Internet.	1-9
2-1	Army hierarchical organization.	2-1
2-2	Total subscriber population.	2-2
2-3	Example force elements, Naval Expeditionary Warfare.	2-6
2-4	1993 view of flag-capable ship long-haul throughput needs.	2-8
2-5	1993 view of flag-capable ship long-haul information categories.	2-8
2-6	JMCOMS shipboard architecture.	2-9
2-7	Connectivity impacts for future applications.	2-11
2-8	Possible future shipboard connectivity implementation.	2-12
2-9	Intra-range communications.	2-17
3-1	Interfaces and boundary of the Warfighter's Internet.	3-3
3-2	Warfighter's Internet network architecture.	3-4
3-3	Providing high impact tactical distributed applications in a resource-constrained environment.	3-6
3-4	Efficient use of uplink and downlink resources.	3-8
3-5	Illustrating the advantage of pooling communications resources (M/M/m model).	3-9
3-6	Uplink and downlink data rates to and from users with omnidirectional antennas.	3-10
3-7	Nominal airborne node-to-individual subscriber downlink performance.	3-11
3-8	Nominal individual subscriber to airborne node uplink performance.	3-12
3-9	Footprint coverage radius (n miles) vs elevation angle (degrees).	3-13
3-10	One foliage scenario.	3-14
3-11	Nominal jamming protection.	3-15
3-12	Terrestrial node links to/from airborne node.	3-16

LIST OF ILLUSTRATIONS (Continued)

Figure No.		Page
3-13	Cross-link antenna patterns.	3-18
3-14	Representative narrow beam cross-link performance (Mbps vs. range).	3-19
3-15	Performance of pancake cross-link antenna patterns (Mbps vs. range).	3-21
3-16	The three ways to connect four nodes with two beams each.	3-22
3-17	Two-step reconfiguration.	3-22
3-18	General architecture of the mobile Warfighter's Internet.	3-24
3-19	Functional view of the mobile WI airborne node communications system.	3-28
3-20	Overview of the Warfighter MCD.	3-32
3-21	Warfighter's Internet and radio access point interconnectivity.	3-36
3-22	Functional view of the mobile WI airborne node communications system.	3-37
3-23	Notional architecture of a representative Warfighter MCD.	3-40
3-24	Base station interfaces to on board communications assets.	3-43
3-25	Architecture (option 2) system level context diagram showing WI components, boundary, and interfaces.	3-44
3-26	Architecture (option 2) is characterized by a Cross-Links Subnet and a Cellular/PCS Subnet, whose nodes and links are shown here.	3-45
3-27	WI airborne node architecture (option #2).	3-46
3-28	Functional diagram of an airborne Mobile Base Station (option #2) supporting Cellular/PCS subscriber links to ground-based Mobile Communications Devices (MCDs).	3-49
3-29	Uplink and downlink virtual channels created by cell multiplexing.	3-50
3-30	MCD and NBC configured as static IP routers for flexibility in connecting other users and subnets to WI.	3-51
3-31	IPv6 address types.	3-57
3-32	Datagram flows to accommodate mobile hosts.	3-58
3-33	Security taxonomy.	3-62
3-34	IPv6 header with optional subheaders.	3-66
3-35	IPv6 Authentication Header (AH).	3-67
3-36	IPv6 Encapsulated Security Protocol subheader.	3-68

LIST OF TABLES

Table No.		Page
2-5	CLZ Control Net (Assault Stage)	2-5

1. OVERVIEW AND EXECUTIVE SUMMARY

1.1 INTRODUCTION

Military operations in the twenty-first century will be conducted in an increasingly information-rich environment. The application of detailed information about the theater of operations (location, composition, and maneuvers of hostile, neutral, and friendly forces) will greatly increase the effectiveness of US forces. To change data into information, sensor data and field reports must flow to analysis centers for processing and then to the users in a form they can digest and apply easily. The increased knowledge provided will allow US decision-making to occur more rapidly than that of the enemy and to immediately take the advantage in the battlefield. Through these mechanisms, information becomes a force multiplier that maximizes the effectiveness of deployed troops. But if information is to do all of this, it must be delivered from anywhere on the globe to the battlefield users, requiring communications systems to deliver this information that reach all force levels throughout the theater. This communications system must also deliver commands to warfighters and allow information transfer and support requests from warfighters back to commanders and support centers. In short, all combatants and their command and support elements must be connected by a secure, robust, global grid of communications as identified in the joint JCS / DDR&E Advanced Battlespace Information System (ABIS) study.

However, this “global grid” today is less than perfect, particularly in the forward areas of the tactical theater. Much of the military communications infrastructure that is currently in the field or being fielded is heavy, large, and not transportable enough to keep up with rapidly maneuvering forces. Its deployment to a rapidly-developing tactical theater is dependent on scarce air and sealift capability. These limitations became evident in many of the recent deployments of US forces overseas. In Operation Desert Storm, the flanking maneuver by US armored forces was accomplished with only minimal communications links because the theater MSE (Mobile Subscriber Equipment) communications system was not designed to keep up with the rapidly advancing armor. In Haiti, US forces landing in Port au Prince had only limited UHF Satcom connectivity with forces landing in the north because of intervening terrain, since heavy SHF Satcom equipment did not land until sometime later.

What is needed is a new communications capability that is lightweight, rapidly deployed, and requires minimal logistic support in the theater. The new capability must be able to provide communications connectivity to isolated forces which are on the move (true mobility). It is also important for it to provide connectivity to the legacy communications equipment found in the Services today and still being fielded. The Warfighter’s Internet concept developed in this report relies on multiple airborne nodes that are within line-of-sight of theater forces, within line of sight of each other, and also connected to command and support centers (some of which are reached through connectivity to the world-wide DISN) (see Figure 1-1). These airborne nodes connected by cross-links form the backbone of a Warfighter’s Internet which provides continuous, robust connectivity to forward theater forces for two-way data, voice, and multimedia traffic. They extend the global networking services utilized pre-deployment into the tactical theater, allowing users to operate in-theater as they are accustomed to in CONUS.

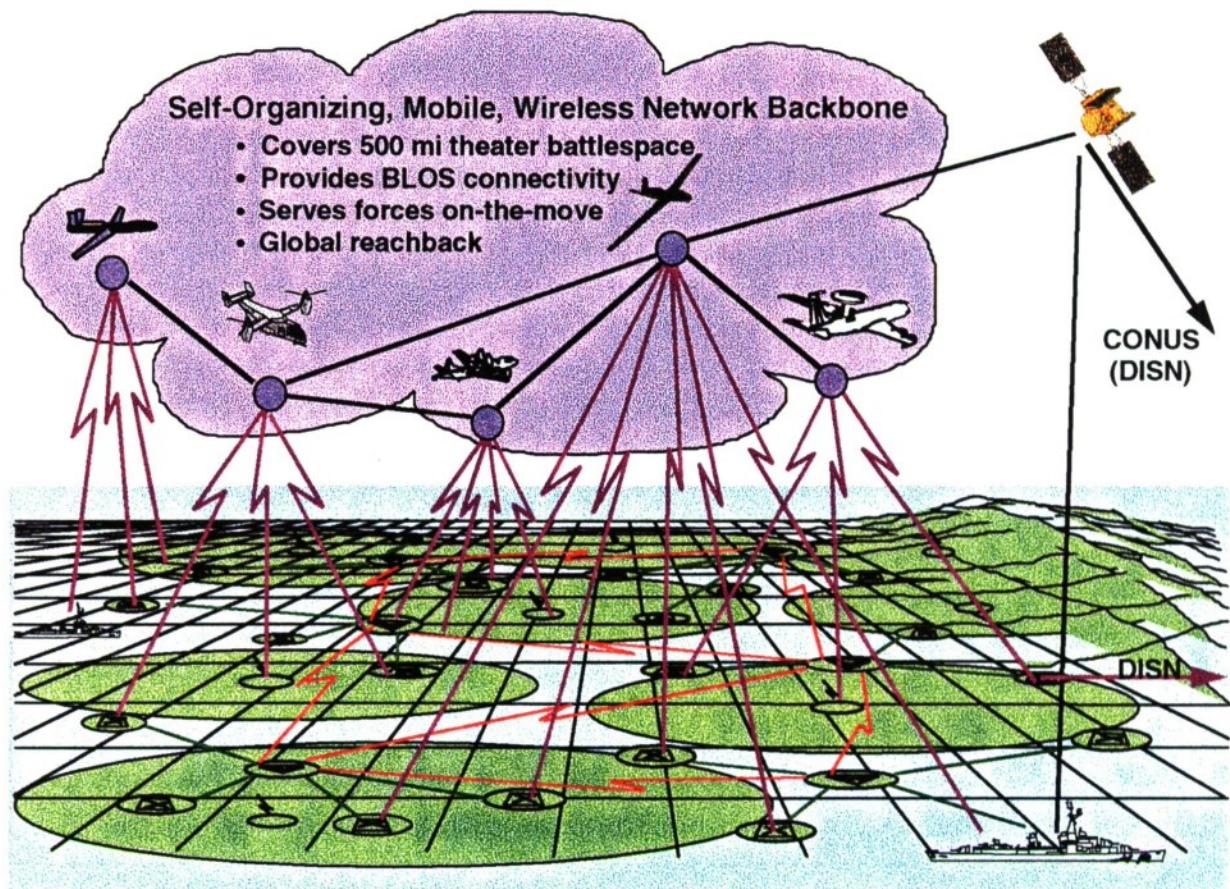


Figure 1-1. Warfighter's Internet concept.

Section 1 of this report continues with an overview of operational scenarios where the Warfighter's Internet will be useful followed by an overview of the architecture and design considerations in the Warfighter's Internet. It closes with a summary of the essential features of the Warfighter's Internet. The report then continues in two more parts. Section 2 details the operational need and concept of use of the Warfighter's Internet with service-specific tactical communications systems in scenarios that may be typical of future contingency actions, as well as its application to T&E and training. Section 3 details alternative technical methods for achieving this capability. Several Appendices expand on specific technical topics.

1.2 OPERATIONAL SCENARIOS AND COMMUNICATIONS NEEDS OVERVIEW

Contingency operations as envisioned in the twenty-first century will be quite different from those of the past. They are likely to be comprised of smaller units, rapidly deployed, and at long range from their support services. Different deployment tactics may be used, such as small unit operations (SUO) where teams of 10 to 20 soldiers are deeply deployed to determine enemy

movements and then call in and coordinate massive indirect fire rather than engage the enemy directly. Special Operations Forces (SOF) today operate in a similar manner, although traditionally they have relied on their own resources for support, and on stealth and surprise for success. The SUO concept, on the other hand, depends on long-range communications (Beyond Line Of Sight - BLOS) for success. However, heavy communications infrastructure is inconsistent with their light armament and deployment methods. The Warfighter's Internet is based on use of simple line-of-sight radios to reach an airborne communication platform which can be used to extend communications to and from Theater and/or CONUS command and support services as an effective way to provide the connectivity these forces need. Through this Warfighter's Internet can flow needed command and control traffic, intelligence, surveillance, and reconnaissance (ISR) products, and requests for medical, logistic, and fire support services. The majority of this traffic is (bursty) computer data, but it is also expected to include (continuous rate) voice and some video (see Figure 1-2).

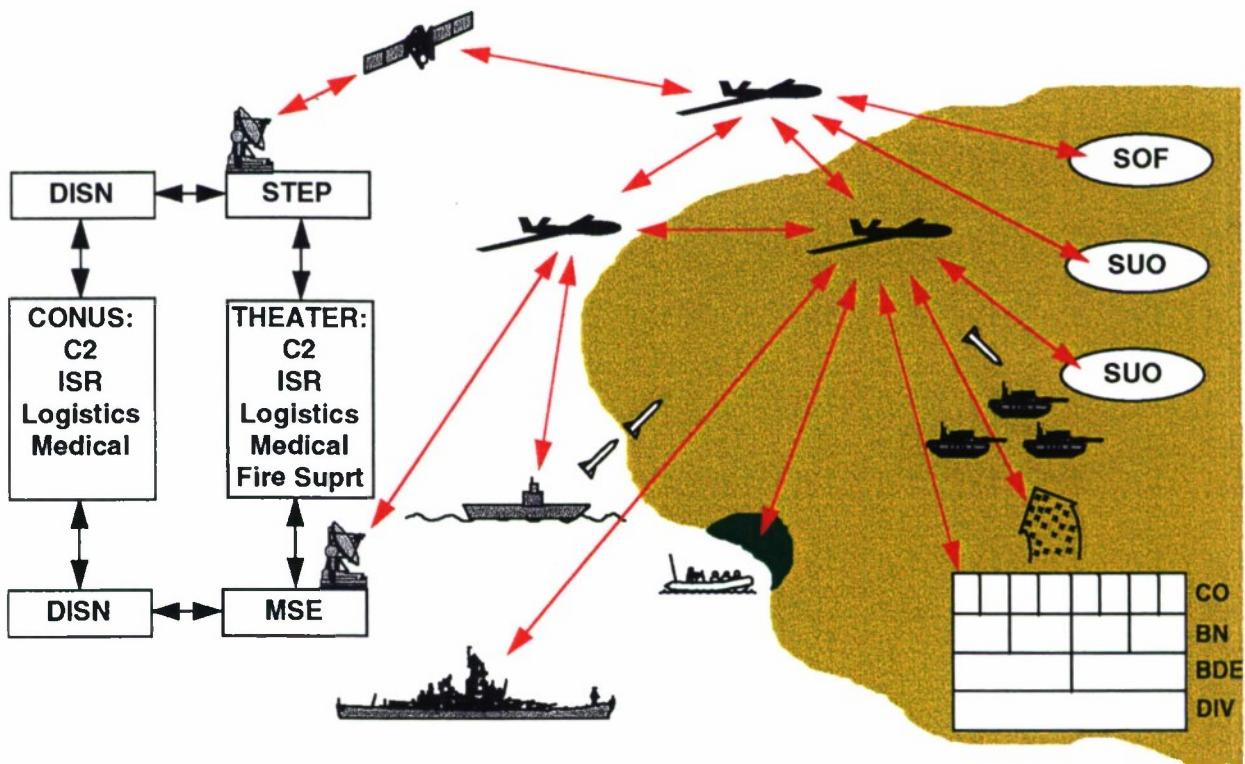


Figure 1-2. Warfighter's Internet provides support of military operations.

Amphibious assault is another scenario that can benefit from an airborne communications network, and is typical of other early entry situations with limited support infrastructure. Future amphibious assaults will likely be carried out with fleet resources positioned over the horizon from the beachhead or landing zone in order to minimize vulnerability of ships to shore batteries and missiles. Troops arriving from landing craft and from helicopters will not have the ability to bring with them heavy communications infrastructure. Again, the Warfighter's Internet will be able to provide the over-the-horizon connectivity needed, while requiring only small, lightweight

line-of-sight radios, some in a handheld or wearable form carried by individual warfighters, and some mounted in vehicles.

Even in more conventional force deployments, there will be situations where some forces have lost line-of-sight connectivity with the rest of the forces, due either to rapid maneuvers or to geographical factors (such as intervening mountains). The airborne network can provide the needed connectivity in these cases as well.

Another way to view the use of the airborne communications assets is as an overlay communication supplement in conventional force deployments. The airborne nodes provide extension of existing theater communications systems to highly mobile or separated forces at or beyond the forward line of troops (FLOT). In this context, it is important that the Warfighter's Internet connect with the lower-echelon ground communications networks found at the edge of battle, such as the Tactical Internet (see Figure 1-3). In addition to servicing these isolated forces, the airborne network can provide alternative connectivity to substitute for deployed MSE equipment which is unable to make a node connection due to terrain, enemy action, mobility, or equipment failure.

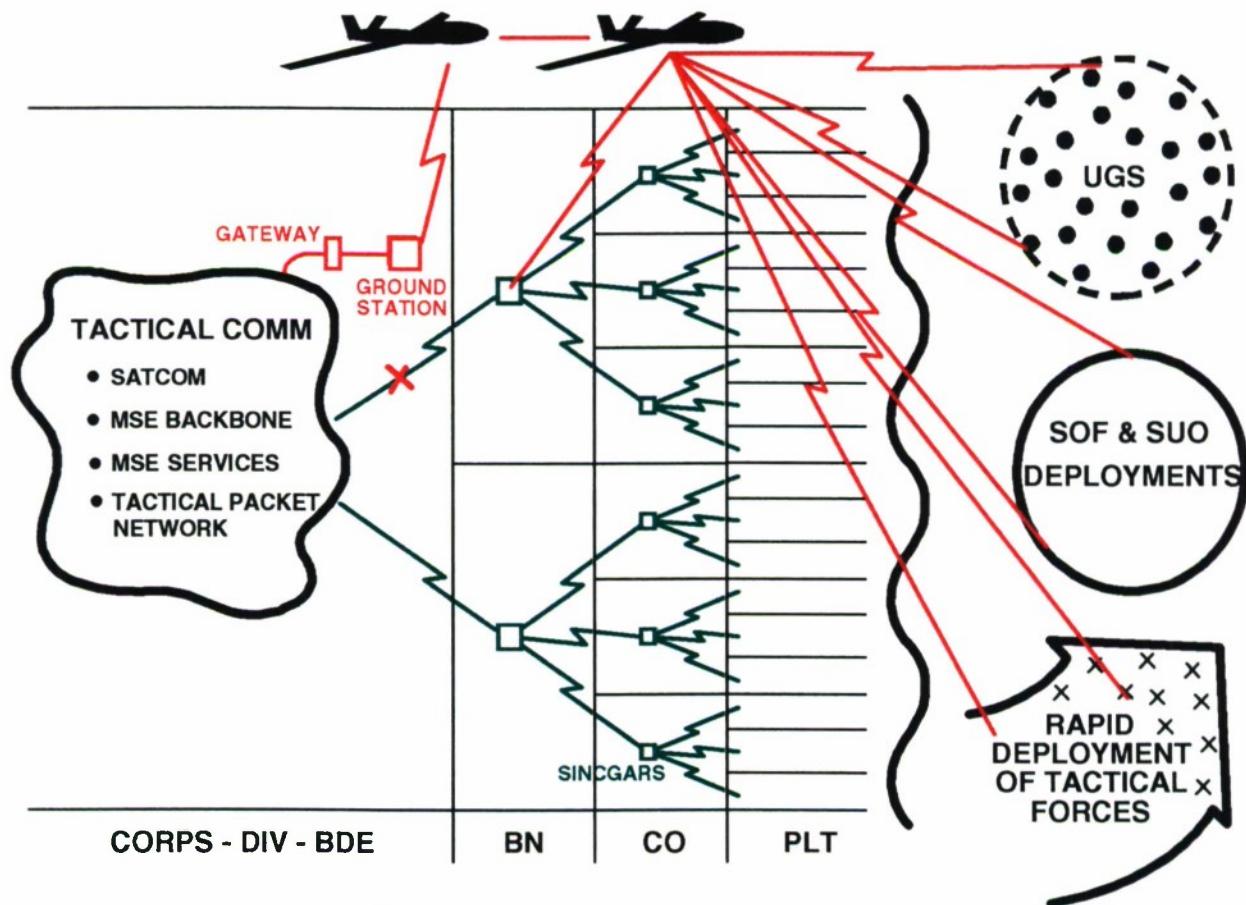


Figure 1-3. Warfighter's Internet as overlay to existing tactical communications.

The Warfighter's Internet contributes in several ways to information flow on the battlefield (see Figure 1-4), and is synergistic with other theater communications systems in the process. Of course the obvious utilization of the Warfighter's Internet is for beyond line-of-sight (BLOS) connectivity for theater-wide networking with global access. This capability would be utilized by Small Unit Operations for C2, support requests, and for ISR data dissemination and requests. ISR products may also be delivered to the theater by the Global Broadcast System/Battlefield Awareness Data Dissemination (GBS/BADD) system to locations equipped with the standard GBS microwave receive terminal and Warfighter's Associate (WFA) processing system. But for forward and lightly-equipped troops, the essential ISR products can arrive through multicast messages over the Warfighter's Internet from the WFA database or through queries of that database by the warfighters. Alternatively, a special T1-rate broadcast might be made (at a frequency around one GHz) of selected BADD information from the ACN airborne nodes to receivers that could operate while on the move using omnidirectional antennas.



Figure 1-4. Integrated information flow in the battlefield.

The airborne nodes used for the Warfighter's Internet include (but are not limited to) Airborne Communications Nodes (ACNs) on Global Hawk HAE UAVs flying at 65,000 ft. The Global Hawk is particularly effective for theater communications because of its altitude (coverage) and endurance. The ACN, in addition to hosting the WI communications equipment, carries other equipment to service legacy radios in the tactical theater (e.g., SINCGARS, LOS UHF, EPLRS, JTIDS). The ACN equipment may also include the T1 rebroadcast of selected BADD traffic and the ability to provide wideband relay between MSE node centers or between RAPs (Radio Access Points, tracked vehicles equipped with the High Capacity Trunk Radio, HCTR). Other WI nodes may also have the capability to connect directly with RAPs. Interconnects between the WI equipment and other ACN equipment can take place on board the ACN platform and/or through the MSE or DISN. Both the ACN and the WI would make use of

the communications satellite terminal on board the Global Hawk to provide global reachback connectivity, so that an in-theater command center would not be required for their deployment.

1.3 ARCHITECTURAL CONCEPT AND DESIGN OVERVIEW

In the communications and computation rich environment of CONUS, military users have grown to embrace the Internet paradigm of information transfer using client-server software, collaborative planning tools, and logistics management programs. Indeed, computer communications has become essential to military operations, which can no longer be accomplished by voice and record communications alone. As the tempo of the battle and dependence on information increases, so will the demand for responsive data communications. The military has recognized this change in operations and has begun to extend an Internet-like data network into the tactical theater. The desire is to have communications capacity and computational resources comparable to that available in garrison in CONUS.

However, at lower echelons (Brigade and below), and before extensive communications equipment can be brought to an emerging theater, communications will necessarily be wireless (not wired) and available capacity will be limited (in comparison with CONUS). The fact that warfighters must carry their communications equipment and operate it while on the move dictates that it be small and low-powered, further restricting the available communications capacity. This, coupled with the need for beyond line-of-sight (BLOS) connectivity (which can be provided by airborne nodes), argues for a very different type of wireless data communications network than has ever been built, and one for which the technology that is needed is not yet available, either from commercial or military sources.

In the first place, the new data network must be very efficient in its utilization of communications capacity to make the limited capacity serve as many users as possible. The fundamental fact that makes this possible is recognition that computer communications is inherently bursty. That is, computer applications send and receive information in bursts (or "packets") that may be separated by substantial periods of silence. Furthermore, subscribers spend more time looking at the screen and typing than in actually transferring data. For a subscriber logged onto a network all day, actual communications utilization may be less than 0.1%. The key to taking advantage of this low average rate is to "statistically share" the communications circuit with many similar subscribers. This is what happens on office LANs (local area networks) and on the wired backbone of the Internet. This is not what happens when a dedicated communications circuit is used. (An example is an individual connecting by a modem to an internet service provider over a dial-up telephone circuit; the circuit is idle most of the time, even though the circuit is being held "open".) Using such statistical multiplexing, communications efficiency (measured in terms of numbers of users that can share the network capacity) can be a factor of 1000 higher than having dedicated circuits for everyone.

The mobile warfighter cannot be given a dedicated or even a "dial-up" (on demand) circuit; there will not be enough to go around. Instead a way must be found to statistically share the limited wireless capacity that can be created between the warfighters and the airborne nodes. One conceptual implementation of such a sharing approach is shown in Figure 1-5. This figure

shows many subscribers transmitting to an airborne node which has several 10s of receivers (but far fewer than the number of subscribers). Because they have bursty computer traffic, subscribers do not each need the full-time use of a receiver in the aircraft. Instead, they request (through an order wire) uplink receiver assignments to send data on a burst basis, and relinquish the channel between bursts for other subscribers to use. For example, 100 subscribers could share one of these uplink receivers if they each transmitted less than 1% of the time, and thus thousands of users could be served by one aircraft. Remember that in this “statistical sharing” model, not all of the 100 subscribers “logged in” to a single receiver will actually be demanding service simultaneously, so at any given instant the capacity of the single receiver is actually being used by far fewer than 100 subscribers; thus it seems to the individual subscriber that he has far more than 1/100 of the full data rate for his use. If the subscriber has a 3-Watt transmitter (peak power), he can transmit to the aircraft at a burst rate up to 64 kbps at a range of 100 miles (with reasonable assumptions on the ground-to-air propagation environment).

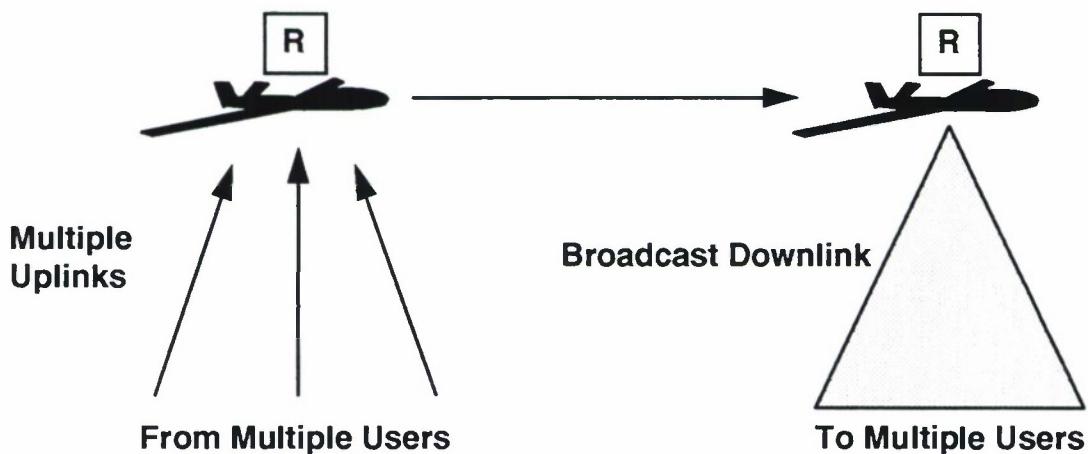


Figure 1-5. Efficient use of uplink and downlink resources.

In the airborne node, the received bursts (or packets) are examined to determine their destination address (contained in the packet header), and are routed appropriately through the network. This routing function (R) requires knowledge of the connectivity of the Warfighter’s Internet backbone (i.e., cross-links) between aircraft and knowledge of to which aircraft the destination subscriber is connected. The packets of data are then placed in a queue along with packets for other subscribers for transmission on the appropriate cross-link. This statistical multiplexing of traffic is an efficient utilization of cross-link capacity for bursty traffic. In fact, it is possible to have quite high-speed links between aircraft; a 25-Watt transmitter using a 9-inch diameter dish can support a 10 Mbps cross-link to a similar dish on another aircraft at a range of nearly 500 miles. Omnidirectional antennas can also be used at lower rates and/or shorter range. Cross-links make it possible to cover very long distances between subscribers while still passing through very few network nodes (minimal latency, simple connectivity), as well as passing over territory that is not yet under friendly control. In contrast, wireless ground-to-ground networks

require individual nodes to be nearly within LOS (about 10 miles apart), and many hops (long latency) with many routers (complex connectivity) are required to cover similar distances.

A similar statistical sharing strategy is utilized on the downlink (shown on another aircraft for clarity). In this case, the aircraft can use a transmitter power that is much larger than that of the individual user, so the downlink data rate can be many times the uplink data rate from one user. For example, an airborne 50-Watt transmitter (operating at a frequency around 1 GHz) can support a downlink data rate of 1.54 Mbps (“T1” in communications nomenclature) to a subscriber with a mobile (whip) antenna. The airborne node assembles all of the data packets to be sent down into a continuous stream; subscribers on the ground listen to the stream for packets that are addressed to them and pull them out. Note that this downlink is particularly efficient for sending the same message to multiple recipients; the message need only be transmitted once with a group address on the packets, and all recipients under that aircraft will hear it. This “multicasting” is a natural for a wireless transmission; multicasting in a wired network (or in a point-to-point wireless network) is much more complicated and less efficient. (See Appendix F.)

The method of statistical sharing multiple uplinks and a single wideband downlink described above has never been implemented before in a commercial or military wireless communication data network. The transmission parts of this system are mostly understood from engineering experience (although significant design choices are still to be made), but the networking implications of this system are revolutionary. New networking algorithms (protocols) must be developed (as evolution and extension of existing Internet protocols) for this system to meet the connectivity and user mobility needs in the presence of mobile airborne nodes and the less-than-ideal propagation conditions in the wireless subscriber links. One challenge is in maintaining the airborne backbone: changing the topology as aircraft move, enter, and leave, as well as quickly updating the routing tables to reflect these changes. These protocol developments are the key technology effort required to implement the Warfighter’s Internet.

The network protocols to be developed for the Warfighter’s Internet must be consistent with evolving Internet protocols, so that software that operates in that environment will also operate in the Warfighter’s Internet. Only this way will it be possible for the military to leverage the vast investment of the commercial world. Some applications will need revision to deal with intermittent connectivity of users (e.g., for reliable multicasting). The WI protocols need to work with evolving Internet Protocols, IPv6, since this will provide some of the needed capabilities such as Quality of Service specification for data, voice, and multimedia integration. (Voice and multimedia must be integrated with data; a separate overlay is not an efficient design.) But additions to IPv6 QoS and multicasting algorithms are needed for variable-quality wireless links. IPv6 protocols also give hooks to allow implementing military prioritization of traffic, security functions, and aids in routing for mobile users. Protocols must also be developed for interfacing with the various signalling and networking methods found in other networks, such as Asynchronous Transfer Mode (ATM), which may be utilized in the backbone of many networks, including the Army’s Area Common User System, Mobile Subscriber Equipment.

An overview of the elements of the Warfighter’s Internet based on this networking architecture is shown in Figure 1-6. The subscriber’s Mobile Communications Device (MCD) is

shown as a hand-held radio (but may be vehicle-mounted) that can be augmented with a personal digital assistant (PDA) or lap-top computer to adapt to data needs and processing requirements, consistent with available power. The Airborne Base Station contains the multiple uplink receivers and the single T1 downlink transmitter, as well as the uplink/downlink resource controller and a traffic routing function. For traffic going to or coming from a cross-link, the Airborne Base Station interfaces to the Cross-link System, which is part of the Warfighter's Internet backbone network. The Cross-link System initializes and maintains the backbone connection topology and routes traffic on and off the backbone. The backbone may also be extended to the surface, providing high capacity connections to ships, command posts, or other multi-user locations. Finally, there is a WI Entry Node, which is where the backbone network interfaces to external networks such as the MSE and DISN data networks. Since this is the primary gateway to the Warfighter's Internet, it also performs a number of housekeeping functions, such as keeping a registry of subscribers and their location, serving as a "care-of" agent to forward traffic from outside to WI subscribers, and functions as a multicast gateway to the outside networks. (Since the Entry Node is vital to interworking with the Warfighter's Internet, it should be redundantly realized; one such realization may be reached over a Satcom link if one of the airborne nodes has a Satcom capability.)

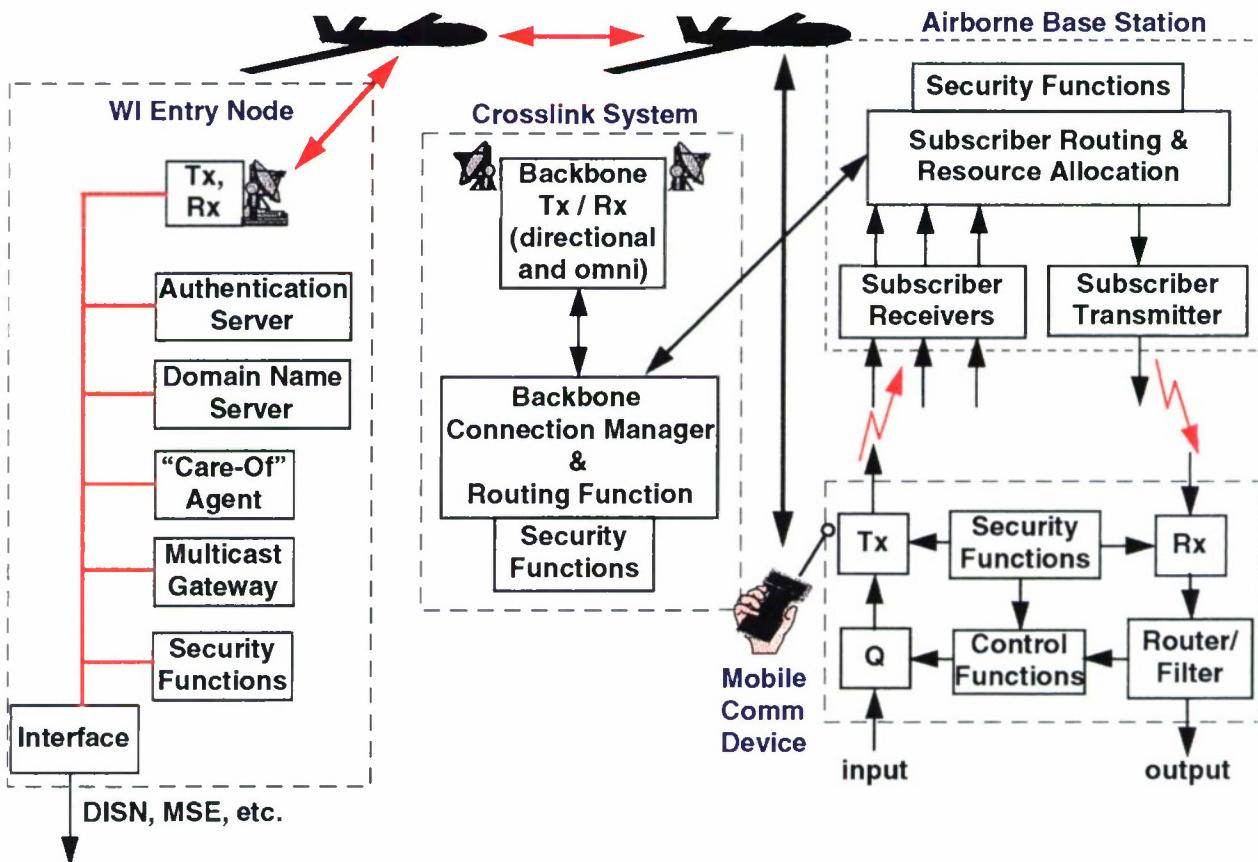


Figure 1-6. Elements of the Warfighter's Internet.

1.4 SUMMARY OF CHARACTERISTICS OF THE WARFIGHTER'S INTERNET

The goal of the Warfighter's Internet is to provide a new communications networking capability for integrated data, voice, and multimedia to on-the-move warfighters in the tactical theater to maintain connectivity to command, support, and information centers. The network utilizes a high-capacity backbone among airborne nodes that serve as basestations for connections directly to subscribers with small, hand-held (or vehicle-mounted) terminals. The backbone can be extended to ground or ship command/support centers, which also provide subscribers worldwide connectivity through other communications networks (e.g., DISN). Because the aircraft can be flown to station from out-of-theater bases, there is no need for a prolonged buildup time to bring in heavy ground communications equipment; thus communications can be operational-ready very quickly. Furthermore, a large coverage area is provided by line-of-sight links to these airborne nodes, allowing service to isolated and highly mobile forces even early in the developing theater operations. This connectivity is needed for anticipated future contingency operations.

The wireless network in the Warfighter's Internet will have a unique implementation of dynamically-shared access to the limited wireless communications capacity so that many subscribers (thousands per aircraft) can be served. This sharing is possible because of the "bursty" nature of computer communications where a subscriber is actually transmitting only a small fraction of the time (perhaps as low as 0.1%), even when logged into the system. (This sharing is similar in effect to the sharing commonly experienced on wired local area networks, LANs.) The end effect is that sharing permits an apparent increase in the throughput to an individual subscriber of a factor of 100 to 1000 compared to dividing the communications capacity up into switched or dedicated circuits assigned to each subscriber. Statistical sharing of this nature is only possible because of the aggregation of capacity in the airborne nodes. Similar gains are not possible with more traditional ground networks (such as multi-hop packet radio, or even commercial cellular/PCS) where each node is modest in capacity and can be shared (at best) by a few users. The simple topology of the backbone nodes in the Warfighter's Internet allows long distance connections to be established that pass through very few nodes. This means relatively simple routing decisions and low-latency delivery of data, and is in contrast to multi-hop ground tactical networks which may require traffic to pass through many nodes (~10 miles per hop) and routers, and where the routing topology is much more complex and dynamic. (Latency is proportional to the number of nodes passed through and the congestion at each node.)

Although the design of individual wireless communications RF links to implement the Warfighter's Internet requires only modest advancements in the engineering state-of-the-art, there are much more significant developments needed in data networking algorithms (protocols) to implement the dynamic sharing of capacity described above and to deal with the mobility of the subscribers and the airborne nodes. Since both the subscribers and the aircraft are in motion, the topology of connections will be constantly changing, and new routing protocols must be developed to respond to these changes in real time. (Today's Internet protocols for the stable, wired environment respond poorly to changes in connectivity; it may take hours for new routing information to propagate throughout the Internet.) The use of wireless links that are relatively unreliable (high error rates; fades) also present new data networking algorithm challenges for

multicasting and for maintaining a specified Quality of Service, since these Internet protocols were developed for (wired) circuits with low error rates and stable connectivity. The development and integration of these new networking algorithms represent the core of the technology challenges to the realization of the Warfighter's Internet.

In summary, the concept of the Warfighter's Internet provides a unique solution to the problem of extending the world-wide information grid to the lower echelons of mobile and/or isolated warfighters. The design concept proposed utilizes statistical sharing of wireless links to serve the bursty traffic of many subscribers. This type of data network has never been built before; it has no parallel in current or anticipated commercial systems (all of which depend on a stable, pre-installed infrastructure). The technical challenge in the Warfighter's Internet is in the development and integration of the new required networking algorithms (protocols) that deal with the sharing of capacity, the mobility of the users and airborne nodes, and the unreliability of the wireless links due to variable propagation conditions. These new protocols must allow the Warfighter's Internet to interoperate for global connectivity with other networks which, along with existing software applications, are based on the evolving protocols of the Internet. By adopting and extending this Internet paradigm for the Warfighter's Internet, it will be possible to leverage from these vast commercial investments in software and networking.

2. OPERATIONAL NEED AND CONCEPT

2.1 OPERATIONAL SCENARIOS AND COMMUNICATIONS NEEDS

2.1.1 Overview

Section 2.1 addresses the utility of the Warfighter's Internet, within the context of individual Service communications systems that are today in use or are planned, for operational scenarios that are likely to be encountered in future contingency operations. The Warfighter's Internet application to Test and Evaluation (T&E) and training is also discussed, particularly for the case of operations in the field on unprepared test or training ranges.

2.1.2 Army Tactical Communications Needs

The Army's primary focus on communication arises out of the operational need to command, coordinate, and control highly mobile major units on the battlefield. The units, their size, and geographic area are summarized in Figure 2-1.

	Element	Commanded By	Typical No. of People	Map Symbol	Approximate Geographic Area FxD Coverage
•	Corps	LTG	200,000 3-5 DIV	XXX	250km x 150km (37,500 sq km)
•	Division	MG	10,000-15,000	XX	50km x 75km (3,750 sq km)
•	Brigade	COL	5,000	X	15km x 10km (150 sq km)
•	Battalion	LTC	800	II	
•	Company	CPT	200	I	
•	Platoon	LT	40		
•	Squad	SGT	10		

Figure 2-1. Army hierarchical organization.

As a result, the Army's tactical communications system for the Digitized Force XXI battlefield has evolved as a combination of systems:

1. At the upper echelon, a circuit switched transportable backbone system (Mobile Subscriber Equipment, MSE) is employed. It is capable of handling large amounts of traffic and is evolving into a wide area ATM system in order to support future multimedia applications.

2. At the lower echelon (Brigade and Below), a highly mobile IP-based network (the Tactical Internet, TI) is employed that handles data and needs to time share data access with voice. This system is currently voice traffic dominated and will remain so for a significant time in the future, although data use is expected to grow significantly due to installation of 1000 computers per Brigade in Force XXI.
3. To enhance the connection between the highly-mobile Tactical Internet and the transportable MSE, a tracked vehicle, the Radio Access Point (RAP) may in the future serve as a concentrator node interfacing to the TI on the one hand, and (through a High Capacity Trunk Radio, HCTR) to an MSE extension node on the other. If this link is BLOS, then the airborne node carrying the WI could also relay this RAP traffic; the WI could then transport this data over its cross-links and/or interconnect with it on board (see Sec. 3.5.2 and Appendix K).

A representative user population is summarized in Figure 2-2.

Area	Subscribers	Total Subscribers
Corps/Div Rear	8000 Voice 1900 Mobile Voice 2000 IP Data	11,900
Brigade	1000 Voice + IP Data All Mobile	20,000
		31,900 Total Subscribers

Figure 2-2. Total subscriber population.

The warfighters and platforms that carry these systems constitute a mobile force that must move rapidly on the battlefield (foot soldier at ~ 3 mph up to vehicles at 60 mph and helicopters at 300 mph). Units can often “leap frog,” move more than once a day, and therefore cannot utilize fixed tower or fixed network resources, but instead must create their supporting mobile communication infrastructure. The ability to rapidly create such a mobile infrastructure is a key difference between military and commercial communications systems and the technology required to implement them.

The WI shares the same high mobility requirements with the lower echelons of the Army Communication infrastructure, but does not, by itself, have to handle the total magnitude of traffic or support all of the theater user population described here, but rather is complementary to standard communications systems (MSE, TI), since it is used to provide connectivity to that (smaller) fraction of the force that cannot be covered by standard systems. The Army’s continued need for real time voice, multimedia, and large file transfer of data in very large volume, particularly at higher echelons, raises serious questions about the capacity of the theater network needed to support all of these services. The WI has assumed its primary traffic will be data from a limited set of users to maintain seamless communications. There will be battlefield scenarios and

services where the WI alone can indeed satisfy these user needs, but for larger deployments and at higher echelons, standard Army communications systems are essential.

There are several scenarios where the WI can provide horizontal and vertical seamless battlefield communications. Three of these scenarios are:

1. Opposed Split-based Lodgment,
2. Opposed Lodgment (ready for breakout), and
3. Rapid Advance-offensive.

In each scenario, the situation involves a limited number of communication users, rapid motion, possible separation from a main force, and the need to eventually transition into the primary communication systems capable of supporting user traffic needs. Eight classes of data traffic have been identified for this subscriber population:

1. Administrative/Logistics: mostly a flow up of traffic, with each command level summing information from below before transmitting further; therefore, messages all essentially the same length regardless of level; currently voice or hand delivery.
2. Intelligence (non-organic sources): largely imagery and textual interpretation flowing into Brigade S2 from higher echelons; only text likely to be propagated below Brigade; flow down limited to affected (and adjacent) units; currently hand delivered.
3. Intelligence (sources organic to Brigade): fundamentally a flow up on a message-by-message basis; includes Contact, Obstacle, Spot, etc., Reports; currently voice; *augmentation by video from FLOT useful in future*.
4. Combat Planning: conference sessions for 15 minutes to 3 hours each, currently largely face-to-face meetings with no relevant data transmissions; *Multimedia capability desired in future*.
5. Orders, including Tactical Overlays: basically a flow down limited to affected (and adjacent) units; currently hand delivered.
6. Friendly Situational Awareness: currently very limited position information, passed by voice; near-future data will flow up from each individual platform and then a union of all (friendly) platforms' positions, etc., will flow down.
7. Enemy Situation reports: real-time distribution of fused information, basically up-dating Enemy Overlay (see No. 5 above); generally a flow down from Brigade and Battalion S2s to affected (and adjacent) units; currently by voice.
8. Coordination: exchanges between adjacent elements and/or combat branches for mutual support (artillery, Medevac, etc.); logical flow mostly between involved units, hence traffic amounts related to level, not numbers of platforms.

1. Opposed Split-base Lodgment involves the first phases of a Brigade-sized unit entering an area to conduct offensive or defensive operations. Under these conditions we would expect heavy reliance on non-organic (i.e., from sources out of theater) intelligence, limited battle planning since the full staff is not yet on site, no extensive orders being issued since no battle planning has been conducted, little situation awareness traffic locally generated since few units are in position, and no need for VTC conferencing with the Brigade Staff or consulting with adjacent units' staff. The majority of the traffic would be coming from remotely located higher echelons in

the form of early intelligence reports and some preliminary administrative/logistics reports. The use of the WI to link to a limited user population is in most cases necessary.

2. During the Opposed Lodgment (Ready for Breakout) scenario, the Brigade is preparing for a breakout. Organic sources of battlefield intelligence coming from forward scout vehicle transmitting video to the Brigade S2 would be involved. Also, non-organic intelligence reports from higher echelons will be used to help gain an accurate picture of the enemy position. If any special forces have been operating in the future battle area, it will be necessary to link this force to the Brigade Staff for effective battle planning and intelligence.

At this time, the issuance or re-issuance of orders will be involved. This can involve the transmission of map overlays, as well as the orders themselves. Combat planning on the part of the Brigade Staff with their subordinate unit commander will be required, and this can be conducted via a VTC.

On a limited basis, situation awareness (SA) messages may be transmitted. This traffic will not be extensive since the unit has not begun to move at this time. The development of a friendly situation picture will be available on map overlays from the SA data. The communication networks associated with this traffic will be a mixture of the Army Tactical Internet, as well as the WI with its airborne nodes. The information from advanced scouts, special force and higher echelon unit can be carried on the WI as a combination of voice, video and data traffic. Local Brigade traffic will use the Army TI; it will be necessary to link the WI into the local network to provide beyond line-of-sight capabilities.

3. Full Offensive, Rapid Advance scenario represents the heaviest load of traffic, both for the Army TI and the WI. The Brigade will be moving out in a Rapid Advance attack. Intelligence reports, both organic and non-organic will be essential. Both the friendly and enemy situation will be changing rapidly. The need for replanning and the issuance of Frag Orders will be essential. Constant updates of the situation via SA data from friendly unit will be involved. The need for unit coordination can require short duration VTCs (~ 15 minutes) with adjacent units. Also complicating the situation will be the loss of connectivity due to geographical obstruction and destruction of radio/computer facilities. The rapid advance of motorized units and armor make continued connection to higher echelons difficult if only the TI is available.

Here, also, a combination of the Army TI and WI will be necessary. The current Army TI is stressed meeting timeline and traffic needs. Additional alternate communication systems will be necessary. Also, assuming that forward elements can send over-the-horizon reports to Division and Corp is important. Augmentation of the current Army backbone system (MSE) that is transportable, but not mobile, with the WI can overcome many of the problems observed in Desert Storm. This further emphasizes the need to successfully interface the TI and the WI. However, the ability to support the volume and mixture of traffic associated with this scenario by the WI alone remains uncertain. Use of Radio Access Points (RAPs) to trunk TI traffic either directly or through airborne platforms to MSE extension nodes may also be needed. This airborne relay capability is being developed at CECOM and is included in the Airborne Comm

Node (ACN) concept; it could also be included in WI airborne nodes, so that the WI could cross-link the RAP traffic or interconnect with it on board (see Sec. 3.5.2 and Appendix K).

2.1.3 Naval Tactical Communications Needs

2.1.3.1 Today's Needs

Today, Naval communications are planned around missions and based on available systems. Planning is conducted prior to deployment, and is based on standard procedures as described in the Naval Warfare Publications and subordinate documents. These documents identify hundreds of communications services and their associated circuits. (A circuit, in Navy parlance, is a particular configuration of equipment. Circuits are described both generically, in a top-level block diagram, and specifically for implementation on a particular ship based on that ship's assets.) The large number of needs identified through this process cannot be supported by available equipment, because these needs are planned for on a one-to-one basis. That is, a search and rescue voice net will require dedicated equipment while it is in operation, even though very lightly loaded. The result is time sharing of equipment such that at any given time, only one of many services is supported by the equipment, and reconfiguration time needed to shift from one service to another is minutes at best, and perhaps hours under some circumstances.

The Naval Research Laboratory led an effort to define communications data needs in amphibious assault, surface fire support, and mine warfare environments supported by SINCGARS data networks. In support of this effort, an ITT report¹ summarized over 100 service requirements for data communications, with Table 2-1 below (from p. 22 of the report) showing a typical small example of the data format. Similarly, the numbers of required voice nets are very large, with multiple nets reserved that are specific to functions within warfare areas.

TABLE 2-1
CLZ Control Net (Assault Stage)

Message Type	Connectivity	Periodicity	Length (bits)
LCAC Position Rpt	LCAC Group CDR to all Nodes	Every 15 Seconds	150
Obstacle Rpt	CLZ CTL TM to all Nodes	1 time in period	1000
SITREP	Nodes to all Nodes	1 time in period	1500

Total traffic assumed in 15 minutes — 14,500 bits.

¹ ITT Industries Aerospace/Communications Division, Joint Littoral Warfare (JLW) Tactical Communications Requirements Analysis Final Report, 31 January 1997.

Figure 2-3 suggests the types of forces and their distribution in a Naval Expeditionary Warfare environment, typical of that used for the NRL-led communications assessment. Features of importance include the wide dispersion of forces, the wide variety of force types, and the wide variety of communications needs implied by the operational and geographic factors. An important factor that may be deduced is the varying need for covertness. Small units operating well forward in enemy territory do not want to be localized, and may not even want their presence to be detectable from electromagnetic emissions. Ships operating well away from land may be less concerned about detectability, but more concerned about potential use of their emissions by enemy radiation-seeking ordnance. Capacity requirements will also vary significantly by unit type, with command elements generally requiring greater capacity to support more services.

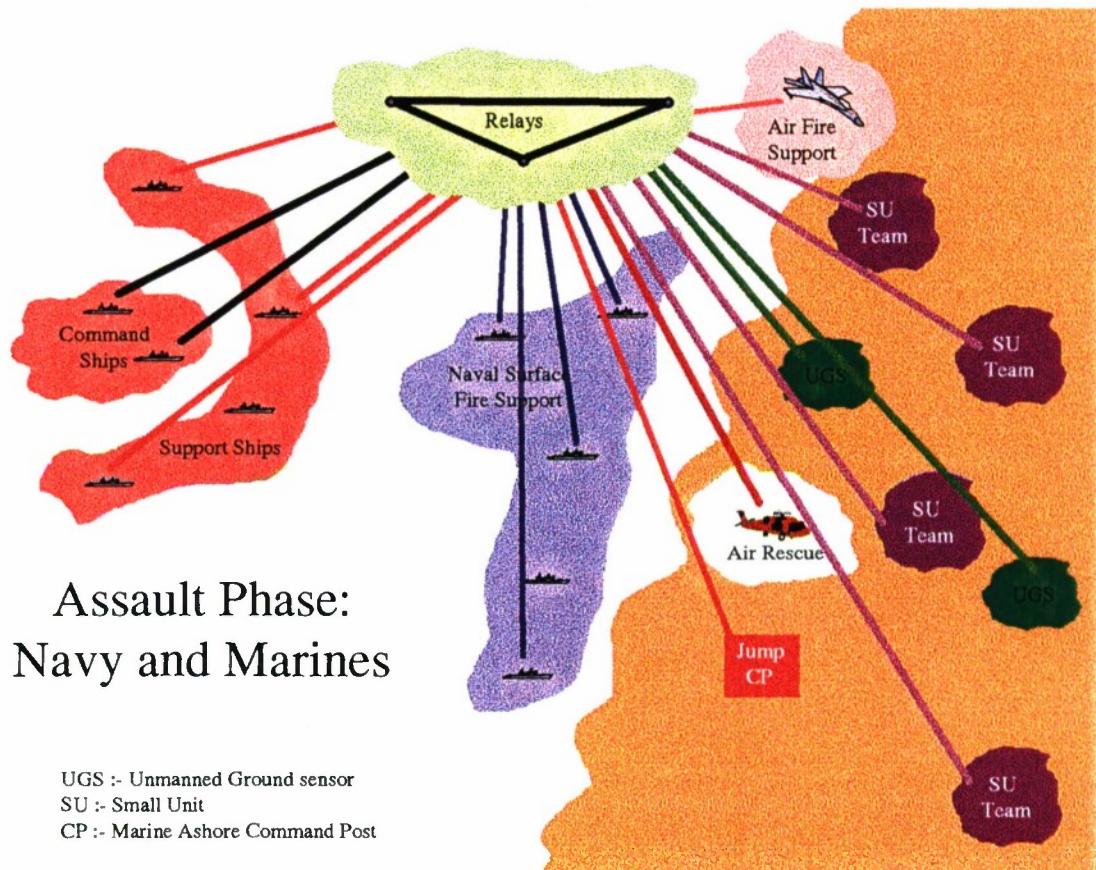


Figure 2-3. Example force elements, Naval Expeditionary Warfare.

The figure suggests a need for connectivity that cannot be supported by line of sight links. Members of forward-deployed teams may operate within line-of-sight of one another, but team-team and team-command connections need over-the-horizon connections. Fire support ships may not be within line-of-sight of spotters today; connections beyond the horizon will be even more common with new extended-range munitions. Minesweeping operations are conducted well away

from other units, requiring beyond-the-horizon connectivity. Both amphibious and carrier group ships will operate 50 nmi or more from shore—well over the horizon. The current practice and continuing trend is to operate most of these ships over the horizon from one another. Virtually all connectivity—for surface fire support, small team operations, isolated unit support (e.g., minesweepers), ship-ship connectivity, ship-shore (both Navy-USMC and USMC-USMC) command connectivity—requires extension over the horizon. This is implied by the relay “cloud” shown at the top of Figure 2-3, which could be airborne platforms. Furthermore, many of these links to landing forces must be to communications equipment that is lightweight and employs simple antenna systems. The only Satcom equipment that meets this description is UHF Satcom, and the capacities of these systems are inadequate for this application. The alternative that is suggested in this report is that the Warfighter’s Internet concept could fill a need that is currently unmet.

But standard doctrine does not call for functions that cannot be met with existing equipment. Video teleconferencing among ships of an amphibious force or carrier group is not spelled out, since equipment to support required data rates does not currently exist on these ships. Yet there is a need for this capability. When a higher capacity research capability (a few hundred kb/sec) radio system was provided to a deploying amphibious group, the equipment was regularly used to support VTC of the group’s “morning briefing.” (In the absence of VTC, this is a physical meeting requiring helicopter movement of key staff members.) In this light, it is unrealistic to expect to find a specific requirement for the Warfighter’s Internet, but the function it could provide is clearly of some utility.

2.1.3.2 Evolving Needs

Growing Throughput Needs

There has always been a need to justify new systems with corresponding requirements, and to address requirements with appropriate systems. Figures 2-4 and 2-5 were prepared by the Space and Naval Warfare Systems Command in 1993 (green elements added to Figure 2-4) to show projected requirements for Satcom connectivity for Fleet flagships and other heavy communications users. Systems were projected to nearly keep pace with demands.

The services called out in Figure 2-5 reflect the capabilities of 1993, but even so lack justification. For example, the figure suggests 100 phone connections at 2.4 kb/sec per connection. Are these numbers sensible? Do these ships always require 100 simultaneous circuits? Is LPC-10 vocoded voice acceptable for all of these circuits, or is better-quality voice really needed (or desired) on at least some of them? Clearly, any assessment of needs is quite subjective.

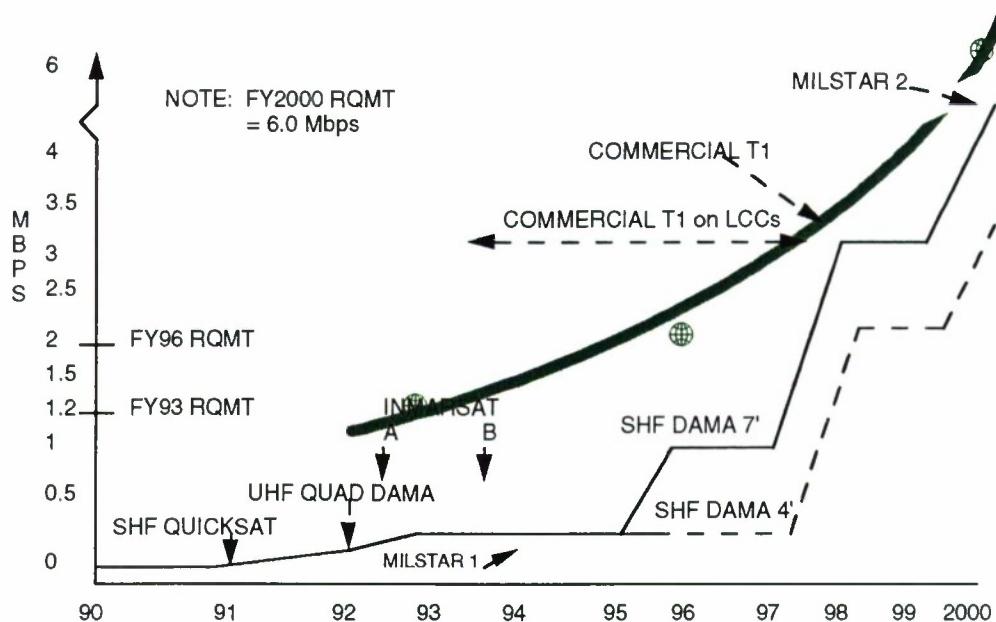


Figure 2-4. 1993 view of flag-capable ship long-haul throughput needs.

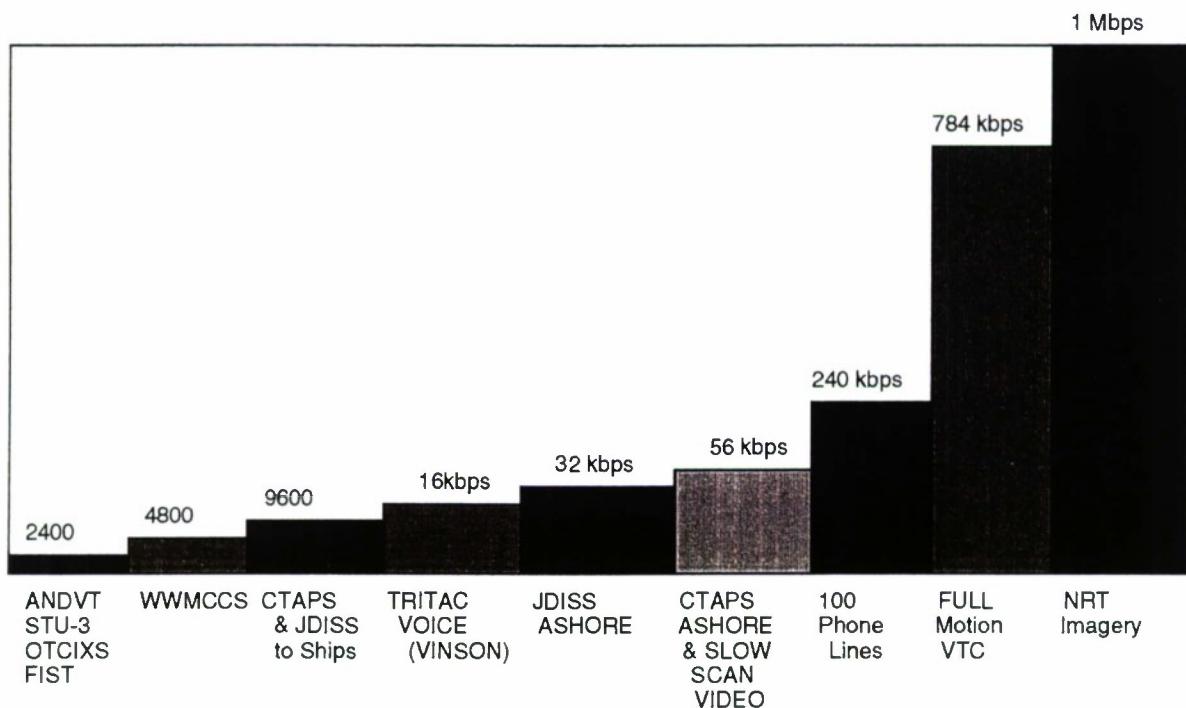


Figure 2-5. 1993 view of flag-capable ship long-haul information categories.

The reality of 1997 is quite different. As a first example of this difference, sensor systems such as UAVs provide high rate “full motion” data streams from visible, IR, SAR, and other sensors downlinked to a single processing location. Once such data exist, multiple users demand access to the raw information—a serious strain for any conceivable affordable ship-ship distribution system. As a second example, the Global Broadcast Service promises data rates to 20 Mb/sec or more, with information sources such as CNN and “quality of life” entertainment video streams, video training, video maintenance, distribution of intelligence imagery, and telemedicine vying for available “space.” The result obeys the law of military communications: “*You build it, we'll fill it.*”

New Systems

Figure 2-6 illustrates the top-level architecture for a shipboard configuration supporting the Joint Maritime Communications Strategy (JMCOMS), and highlights aspects of off-ship access. For aircraft and submarines, a different set of media would be shown, while shore facilities (other than communications facilities) would generally have only landline external connections. For any node, there are four major elements of the architecture: users, switching and routing for interconnecting users either locally or via “off-node” connectivity, various communications media for providing wide-area connectivity, and connectivity management.

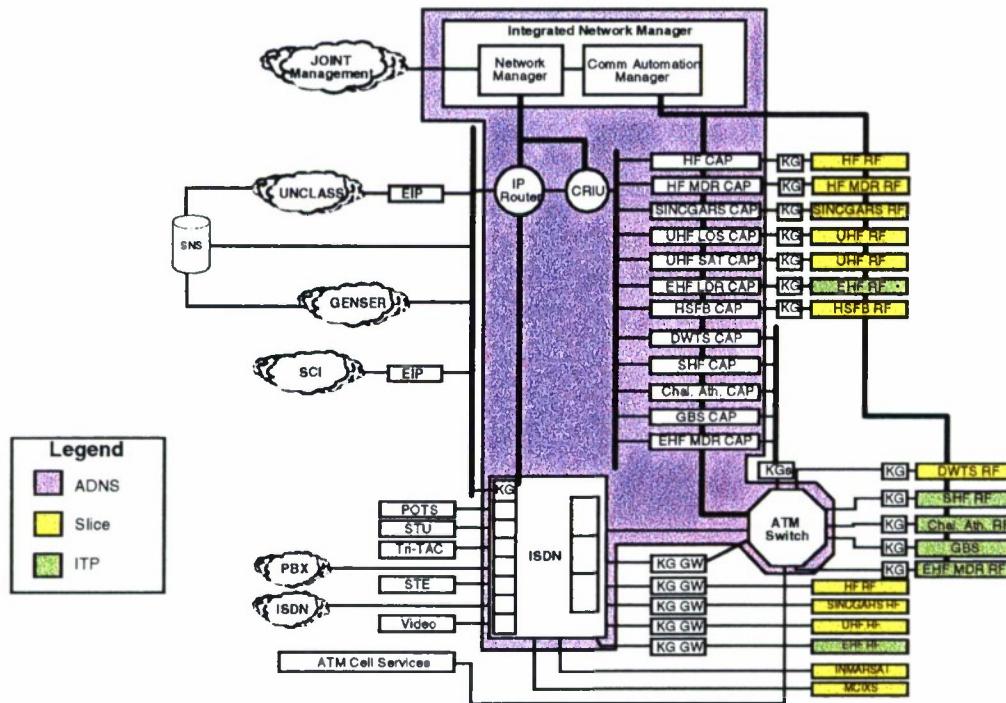


Figure 2-6. JMCOMS shipboard architecture.

JMCOMS is intended to provide highly flexible and responsive connectivity, in addition to its major goal of sharing available resources among all users in order to provide high quality service to support the growing demands of more users. On the figure, the elements identified as "CAP" (communications access protocol) serve as interfaces between the switching and routing system and the individual radio systems. The implications for Navy ships of adding Warfighter's Internet access might involve implementation of a CAP for WI, or possibly direct ATM access similar to that envisioned for high capacity systems such as SHF Satcom. The net result is that suitable radio equipment (such as the WI backbone radios) could be added to any ship, with access by all shipboard users, and any required gateway functions, easily provided by JMCOMS.

New Services

User access and formats for the military will follow and exploit commercial technologies. The current commercial trend has been towards greater data rates at all levels from user equipment to the public switched networks. Over the past two decades, each of these rates has been growing steadily at about 38% per year.² This growth on the user side has been spurred by both technology and user demand. On the switched network side it has been driven by gradual increases in aggregate user demand, rather than by dramatic increases in the demands for high capacity. By far the largest consumption of the total public switched network capacity today is by individual voice users.

Much of the high visibility discussion of the future high capacity public network focuses on consumer applications such as movies on demand, interactive video games, and home shopping. Closer to Navy interests are the trends in business applications driving high capacity connectivity requirements. Figure 2-7, from Sakata,³ shows applications on a delay sensitivity versus data volume grid, indicating the gigabit domain. Note that many of these areas, particularly groupware applications, multimedia teleconferencing, remote visualization, and virtual reality, may be highly significant areas for future Navy applications. Of course, for mobile nodes, applications in the gigabit domain will be confined internally to the node itself due to lack of gigabit external connectivity, except possibly pierside.

²Ransom, MN, and Spears, DR, *Applications of Public Gigabit Networks*, IEEE Network, p. 30, March 1992.

³Sakata, S, *B-ISDN Multimedia Workstation Architecture*, IEEE Communications Magazine, p. 65, August 1993.

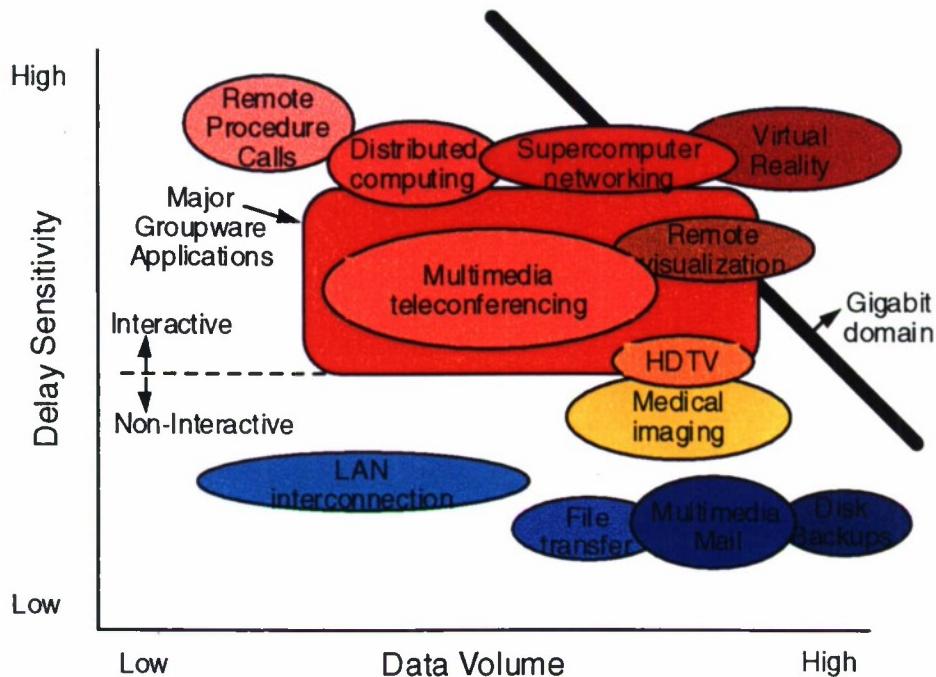


Figure 2-7. Connectivity impacts for future applications.

2.1.3.3 Goal Properties

Figure 2-8 suggests a future shipboard physical implementation of the architecture, based on Asynchronous Transfer Mode technology. The multiple shipboard networks currently required to support computer (local area network) and voice connectivity are all replaced by a single, redundant, fiber-optic ATM-based network. Storage devices for various real-time and library databases, and sophisticated special-purpose processors are connected by the network. User workstations and voice terminals all directly access the ATM network at the standard rate of 155.52 Mb/sec via optical fiber. Future special-purpose systems such as 3-dimensional displays or virtual reality applications can also be supported by the network.

Off-ship access from any workstation is available directly on some links supporting ATM modes (pierside, high data rate RF). Off-ship access is provided via conventional (non-ATM) media through gateways connected to the shipboard ATM network.

While terminals, such as the multimedia workstations shown on the figure, are not directly connectivity elements, their underlying architecture can impact connectivity requirements. The Navy is moving away from UNIX systems which dominated shipboard installations (the Navy Tactical Advanced Computer programs, TAC-1 through TAC-4, are UNIX machines), in favor of less expensive PC and Windows environments. The Navy will likely follow industry in any move away from the PC and towards even less expensive, but highly capable, network machines.⁴ Such

⁴ Comerford, R, The battle for the desktop, IEEE Spectrum, p. 21, May 1997.

a move would be in the direction, envisioned over the past several years, of the user-defined workstation which autoconfigures to meet the needs of the user, rather than being dedicated to a single function such as anti-submarine warfare. (Dedicated workstations and work spaces are the case today.) The implications are a need for high capacity connectivity to servers for initial configuration or reconfiguration, including downloading of software and current data sets—as shown in Figure 2-8.

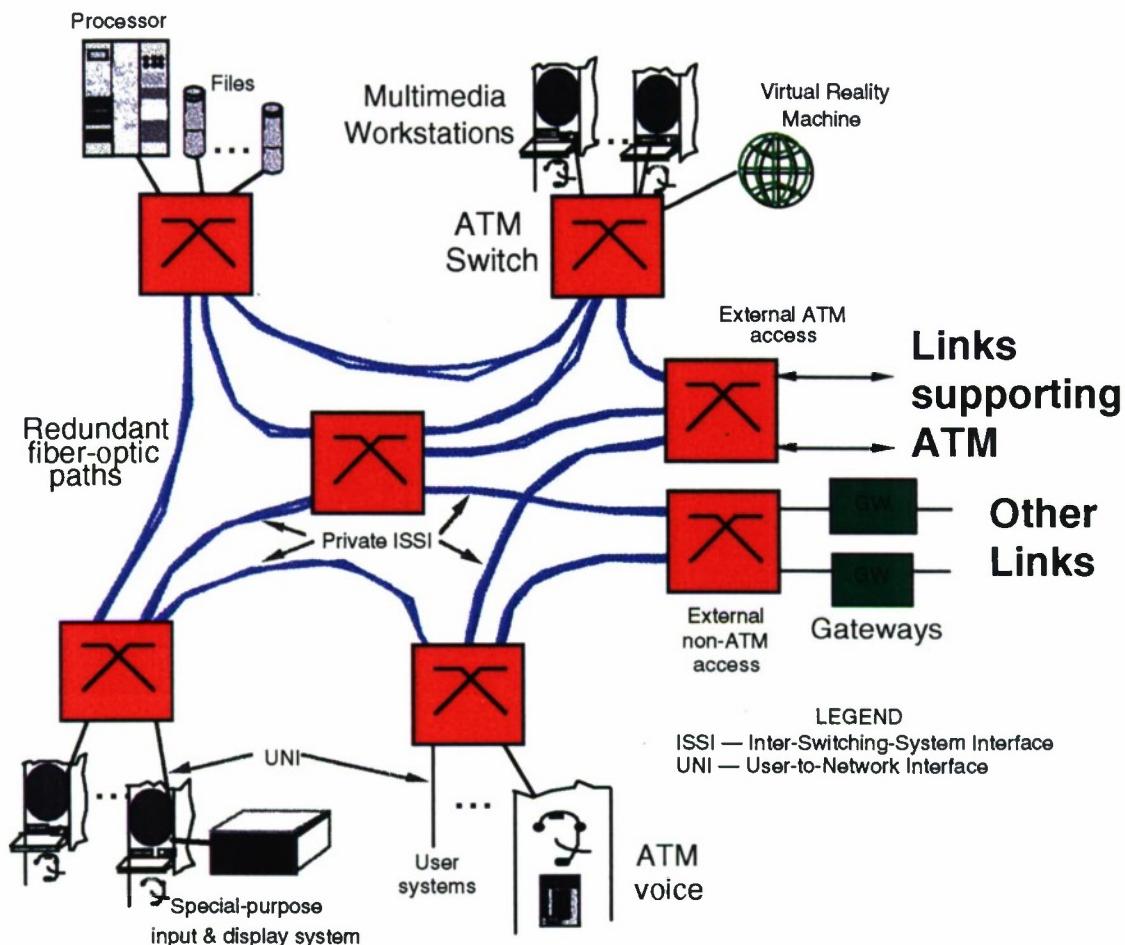


Figure 2-8. Possible future shipboard connectivity implementation.

Key attributes of future communications include the following:

- **Connectivity (endpoints).** The system must support rapid user access to all other relevant users and information sources required to carry out missions and functions. Such connectivity must be dynamic, adapting to changes in real-time needs of the user. Users must be able to easily establish single-user or group connections, adding or deleting members as needed.

- Information format. The system must allow the user to define an appropriate format for the business of the moment, and to support rapid changes in format as dictated by the situation. “Format” here implies the traditional voice, data, video, imagery, conferencing, and emerging formats such as “whiteboarding” and virtual reality.
- Quality of Service. The user must be able to define the required quality of service, including error tolerance, timeliness, priority, security level, precision (including, for example, voice quality or video distortions), and vulnerability to intercept, direction finding, or exploitation.
- Throughput. The user should be able to define desired or required traffic throughput rates, or volume of information to be transferred and allowed delivery time. This factor is directly tied to quality of service.
- System responsiveness. The system should respond to “reconfiguration” requests by the user (for example, opening new group “connections”) in real time.

2.1.3.4 Warfighter’s Internet Implications

The needs discussed may be summarized in the following points:

- Force dispersion and low densities imply a need for over-the-horizon connectivity.
- Throughput demands are growing based on “new” types of services. Ship-ship and command ship-shore needs can be expected to grow to fill whatever channels that might be provided using current technologies.
- Shipboard services will be carried via commercial network technologies, including IP and ATM.
- The advent of JMCOMS makes the introduction of new link types simple, so long as the links support conventional formats such as IP and ATM.
- Small, mobile forces may not have the logistic support to employ other than small, lightweight communications equipment with simple antenna systems.

High altitude relay spanning the littoral theater, such as Warfighter’s Internet or Satcom systems, are clearly required to meet these needs. For the foreseeable future, worldwide coverage, cost, security, and capacity needs will not be met entirely by Satcom systems, and the cost / size / logistics of Satcom terminals may not be acceptable to small, mobile forces. Thus the concept of the Warfighter’s Internet offers a solution to a real problem.

2.1.4 Air Force Tactical Communications Needs

The Air Force has operational communications needs both global and tactical. In this section, the Air Force communications needs in the tactical theater are described, and the potential for implementation of some of them through the Warfighter's Internet is identified.

In a Major Regional Conflict (MRC), the Air Force will deploy significant assets into the tactical theater to support and control air operations. Included in these assets are the theater Air Operations Center (AOC), multiple Tactical Air Bases (TAB), the Theater Air Control System (TACS) [consisting of the Control and Reporting Centers (CRC) and Elements (CRE) (radars and comm links to aircraft), the Air Support Operations Center (ASOC) and Tactical Air Control Party (TACP) (which coordinate close air support with ground troops), and the airborne command, control, and sensor elements AWACS, ABCCC, and JSTARS] and the Tactical Airlift Control Element(s) (TALC) to control airlift operations. In a Limited Regional Conflict (LRC), the number of these elements will be greatly reduced, and many of them may not be positioned in-theater, but in nearby friendly countries with significant support infrastructure. For an MRC, communications to these widely deployed assets is a major concern; for the LRC, theater communications requirements are substantially lower. (However, communications with TACPs will remain vital to coordinate close air support.)

The Air Force has plans to support most of these ground elements with satellite communications systems that make use of available Milsatcom resources such as DSCS, UFO, and MILSTAR. Alternatively, if the elements are co-located with other US forces, communications assets may also be drawn from, for example, the Army ACUS and MSE. Satellite communications is particularly suited to the Air Force needs, since many of the Air Force elements are widely dispersed. However, in early entry when limited communications satellite terminals have been brought into theater, the Warfighter's Internet can serve as a substitute for the satellite communications. That is, the WI can provide coverage of large areas and therefore connectivity of isolated elements that each connect line-of-sight to an airborne node of the WI. The Air Force user on the ground would need only the small, lightweight WI Mobile Communication Device, supplemented by a notebook computer. This capability would also be particularly useful for the TACP to provide connectivity to the ASOC and AOC.

Communications to aircraft through the WI are also possible. Today a TACP communicates to strike aircraft via LOS radio. Because strike aircraft prefer to hug the terrain en route to the target, there is a very limited time when communications to them is possible from the TACP. Alternatively, the WI (high above the battlefield) can relay communications between the incoming strike aircraft and the TACP over a wide range, as well as providing reliable connectivity between the TACP and the ASOC. Likewise, the WI can deliver target and threat updates to the cockpit of the strike aircraft over LOS links. The WI nodes could relay information to/from airborne control and sensor platforms (AWACS, ABCCC, and JSTARS) from/to other theater locations as well.

2.1.5 Warfighter's Internet Applications in T&E and Training

2.1.5.1 Introduction

Training of military personnel is a non-combat function that is recognized as being critically related to the success of military operations. Training exercises are currently carried out on special instrumented ranges which allow monitoring of the exercise and evaluation of unit performance in simulated combat conditions against a simulated enemy. Three important trends in training must be considered in the future: 1) use of virtual forces (simulated forces, including man-in-the-loop simulators) coupled to the live range exercise, 2) more joint service training exercises, and 3) conducting more exercises off-range, i.e., on unprepared sites/areas to achieve greater variety and realism in mission rehearsal.

Exactly the same trends can be seen in future test and evaluation (T&E) activities which take place throughout the system acquisition process, and continue after actual system deployment. Development test and evaluation (DT&E) is usually done on a small scale to collect engineering data on performance of the system under test (SUT), requiring telemetry that will allow engineering diagnosis of malfunctions of the SUT. Operational test and eval (OT&E) can be done on a larger scale, perhaps on the order of a significant training exercise, but the data collected is oriented toward evaluating the accuracy or ability of the SUT to cope with stressing scenarios. Operational tests are usually done with military personnel operating the SUT.

Both DT&E and OT&E can include virtual forces to enhance the scope of the testing, for example, by overloading the SUT with virtual targets or ECM threats. The continued OT&E after deployment is concerned with T&E of system upgrades, performance against new threats or scenarios, or new employment concepts. Both DT&E and OT&E are expected to take place off-range more frequently in the future, requiring special communication support to couple the virtual entities into the live tests and to collect and analyze the data from a remote area. This is where the Warfighter's Internet appears to be of the most utility in T&E.

2.1.5.2 Test & Evaluation (T&E)

DT&E is involved in the system development process while OT&E activities address the effectiveness of systems, doctrine, and operational procedures. In general, however, a test involves one or more test ranges, a range control center (RCC) for each participating range, and (possibly) a virtual battlespace. The virtual battlespace includes remotely located simulation facilities that must be coupled to the players and systems on the live range, as well as its own control center, which we will refer to as the "virtual battlespace control center" (VCC). The test range has a number of components:

- Live entities operating the equipment or carrying out a particular maneuver,
- A variety of platforms including the system under test (SUT), the focus of the test,
- Targets or opposing forces to challenge the SUT,
- An instrumentation system to locate and determine the state of the entities,
- A telemetry system for gathering data on the SUT operation,

- A communication infrastructure to collect data, monitor and control the test, and
- A data logging facility to gather and archive data from the test.

When more than one test range or a remote virtual battlespace is involved, the communication infrastructure must be augmented to tie these facilities together so they can interact properly during a test. In this case, the RCC, whose prime responsibility is the safety and control of activities on its live range, must also be augmented with a higher level control center for overall control of the multiple facility test. We refer to this entity here as the “overall test control center,” or the OTCC. The OTCC may be collocated with one of the RCCs or the VCC. These facilities are called out because of the implications that they raise with respect to the communication infrastructure.

There is a similarity between T&E activities and military training exercises, but in T&E greater emphasis is given to collecting data on the system under test (SUT). This even includes provision of real targets (remotely controlled) for *live fire* testing. There may be a single entity SUT (such as a Patriot air defense fire unit or artillery locating sensor) or many entities of the SUT (as in the case of a new rifle or anti-tank weapon).

T&E Scenarios

Currently most T&E activities take place on prepared test ranges with live participants operating the equipment to be evaluated. The area required for DT&E may be quite small, depending on the SUT requirements. Testing a new weapon with 1 km range would require on the order of a few km. However, in medium and long range missile tests, two separate ranges are often involved, and communication and coordination between them is complex.

OT&E could resemble a large training exercise, with a large number of personnel organized into Blue (friendly) forces and opposing forces (OPFORs). Tests may involve up to a battalion-sized unit, which, together with a battalion-sized OPFORs, would require on the order of 1000 sq km, with a separation between entities of as much as 100 km. Current test ranges usually provide a communication, control, and instrumentation infrastructure. New types of tests, perhaps involving joint forces for the first time, may require new or augmented comm facilities.

Due to limited choices of environmental conditions that are available at current test ranges, it is often desirable to conduct certain tests off-range, i.e., at unprepared sites. If a tropical rain forest environment is needed for T&E of a new soldier radio, present test ranges may not be capable of providing adequate realism for such environmental conditions. Similarly, systems already deployed on operational platforms are often not able to go to a test range, and a system upgrade might have to be evaluated at an operational military base or even in the open ocean. In the case of off-range T&E, there is one set of comm problems at the test site (i.e., intra-range), and one associated with comm requirements with remote facilities that could be thousands of km away (inter-range).

Participants, Types, and Numbers

DT&E involving a single, large SUT might require high rate telemetry from the SUT (on the order of tens of Mbits/sec) and perhaps ten or more remotely-controlled targets (each requiring a control link providing a data rate of hundreds of bits/sec). Personnel counts might reach up to 100 people, all of whom should be monitored (position and state or condition) while on the live range. The SUT may be faced with a number of live targets, remotely-controlled targets, and virtual targets (simulated entities remotely located, including human-in-the-loop simulators).

At the other extreme, OT&E might involve a battalion vs. battalion simulated battle, including several tens of tanks, Bradleys, rotary wing aircraft, and perhaps even fixed wing aircraft. The typical army battalion has 600-800 members. Two battalions plus perhaps 100 observer/controllers (O/Cs) on the range to monitor and control the activities, plus other supporting units, could result in nearly 2,000 live entities on the range during such a test. SUT telemetry would still be desired, although lower data rates would suffice, since OT&E requires data sufficient to determine exactly what the SUT did, i.e., how well it performed, but not for diagnosing engineering design problems.

Comm Services Required

On a single test range, for a simple test on a SUT, the basic types of intra-range message flows can be summarized as follows (see Figure 2-9):

1. Data reporting (position and status) of live entities on the range to the RCC,
2. Control messages to the O/Cs, remotely controlled devices and platforms, and perhaps to the live players (e.g., “stand by - test on hold,” or “break it off”), and
3. SUT telemetry and other special instrumentation sensor data reporting to the RCC.

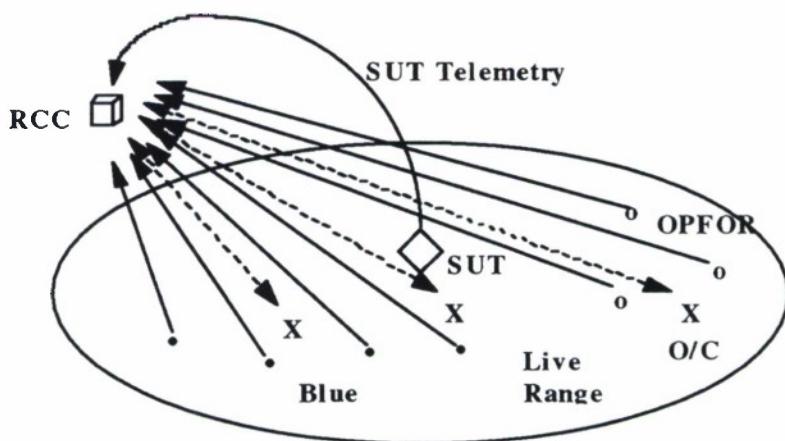


Figure 2-9. Intra-range communications.

In the simplest case of a single range and a single RCC, all of these message types can be handled most conveniently with a star network topology. This is feasible, for example, if LOS links can be achieved between all players on the range and the antenna tower(s) at the RCC. If LOS cannot be achieved (because of the length of some links or terrain blockage), then comm relays are necessary, and the topology becomes more complex, as does the net management and control.

Data reporting to the RCC could, in general, be satisfied with a connectionless messaging service. The latency associated with a wireless network with transmission lengths no longer than 100 km should be on the order of milliseconds (depending on data rates and loading, of course), which is more than adequate for this function.

Data reporting by a live player on the range involves sending geolocation and time of measurement (ranging from 64 to 256 bits), and status (8 to 16 bits) in a message to the RCC. The update rate depends on the speed and acceleration of the entity. For example, a tank would probably require no more than one update per second. The more dynamic entities such as fixed wing aircraft and missiles would require up to perhaps 10 to 30 updates per second, while the dismounted players might require only one or two updates per *minute*. A very rough upper bound on the position and status reporting message length to the RCC can be taken as 256 bits before error correcting coding. Rate one half error control coding (ECC) is assumed, so that the total coded message length for this service would be approximately 500 bits in round numbers.

A large scale test involving two battalions and 100 O/Cs would require an aggregate communication capacity of approximately 35 to 50 kb/s for the coded data (i.e., including the rate one half ECC redundancy). A total of 25 aircraft with update rates of 10 per second would add another 125 kb/s. The total capacity required for entity reporting to the RCC may be upper bounded by 175 kb/s.

Another form of monitoring the live range involves the transmission of imagery or video from a number of the O/Cs or fixed camera sites to the RCC. Assuming a 1000 x 1000 pixel (still) image, 24-bit color, and 10:1 compression, a single image would require the transmission of 2.4 Mbits. A high-speed data service operating at a rate of 120 kb/s could transmit a single compressed image in 20 seconds (the suggestion of 120 kb/s is motivated by wireless link considerations). This could provide a maximum of 3 independent images per minute to the RCC, which is minimal. A large-scale exercise might require 5 to 10 times this image transmission rate. Video would be even more demanding than compressed still images, especially if full motion, TV resolution video were required. This falls into the SUT telemetry category.

Control messages from the RCC to entities on the range, shown in Figure 2-9 as dashed lines, would use both data messages and voice (discussions between the RCC and O/Cs on the range). These messages would be sent almost exclusively to the O/Cs and very rarely to a subset of the players. Compared to the capacity required for entity reporting, the capacity of these control messages is dominated by the number of voice channels that are simultaneously in use during a test. Assuming 20% of the O/Cs use voice during a test, and 4.8 kb/s digitized voice (full

duplex), this amounts to a total of 200 kb/s. Latency requirements are also governed by the voice channel requirements of 100 to 150 msec for acceptable two-way conversations.

Telemetry from the SUT can vary widely, depending on the type of test and the type of SUT. In DT&E, the telemetry would be more demanding than the other services simply because the scale of the test is likely to be smaller and the telemetry is more critical. In OT&E, a large scale test would involve a large number of players on the range but less telemetry from the SUT. If we assume SUT telemetry for DT&E requires 100 times the message length as entity report messages and 100 updates/sec, this implies a telemetry link with a capacity of 5 Mb/s. The telemetry can clearly dominate the communication requirements for T&E. However, the requirement is for connectivity only between the SUT and the RCC. A point-to-point telemetry link would suffice. If a more complex SUT requires significantly higher data rate than the 5 Mb/s estimate, then a separate point-to-point link may well be the cost-effective solution.

2.1.5.3 Training

Military training has as its objectives the development of a military force that will be very effective in all aspects of military combat, and maintaining the readiness of this force. Training activities range from initial basic training of recruits, to specialist training with military equipment, procedures, and tactics, to training exercises to bring military units to peak effectiveness. Mission rehearsal in preparation for an expected deployment is an additional training activity. One of the keys to military success has always been to put the best equipment in the hands of well-trained personnel. In view of the demands being placed on the military in the 21st century to carry out a greater variety of missions with reduced forces and more complex equipment than ever before, the training issue assumes even greater importance than in the past.

The focus of this discussion is on the more advanced levels of training, i.e., training exercises and mission rehearsals which take place on training ranges. The reason is that training ranges have unique communication needs to which the Warfighter's Internet may be particularly well suited. These ranges may be located at training centers, but in the future they may frequently be unprepared sites located almost anywhere in the world. As with continued T&E activities after system deployment, training must be a continual process even for deployed forces, and ability to use unprepared sites in the vicinity of the deployed units would be logistically very convenient. The use of unprepared sites for mission rehearsal also allows more flexibility in choosing locations that can closely replicate the expected environmental conditions of an upcoming deployment or a region that may be of interest in the future.

Training exercises may be quite similar to T&E activities such as operational testing, but while the T&E focus is on the evaluation of the SUT, the focus of training is always on evaluation of the effectiveness of the military unit under prescribed conditions. In some exercises, both T&E and training objectives may be addressed. One example of this is the series of annual exercises conducted by the All-Services Combat Identification Evaluation Team (ASCIET) which serves as a vehicle for training of joint forces as well as evaluation of target identification techniques and equipment.

Conducting Training Exercises

Training exercises, like T&E activities, must be monitored and controlled while in progress for reasons of safety and adherence to the exercise objectives and plan. Data must also be collected and stored or logged to allow analysis and replay of the exercise (as a training tool) and evaluation of player performance. Additional functions required in training exercises but not in T&E exercises are scoring of combat engagements (deciding on the outcome of shots or missiles fired), and kill notification and removal without disruption of the exercise.

A training range must provide a realistic and ample physical environment within which training exercises can be conducted safely and efficiently. Blue forces (the trainees) must maneuver, move to contact, and engage the opposing forces (OPFORs) within the bounds of the range. Some training exercises in land warfare may be conducted in a geographic area as small as 10 x 10 km. However, certain exercises involving rapid maneuvers may require movement of units between locations separated by as much as 100 km. Aircraft participating in an exercise may use remotely located bases and fly to/from the combat area (the training range), or they may be based at air strips within the range.

A training exercise may be conducted on a single range, or may involve coordinated maneuvers of forces on multiple ranges. Live exercises can be augmented by virtual forces (simulated entities located on a virtual battlespace) to expand and enhance the training experience. As mentioned earlier, some exercises may involve the use of an unprepared site as a live range. All of these levels of training exercises deserve consideration as potential applications that could be usefully served by the Warfighter's Internet.

Single Range Training Exercise

Consider first a training exercise involving a single live range. Such an exercise is usually managed from a Range Control Center (RCC) located at the edge of the range. The functions of the RCC are:

- Monitoring the state of all entities on the range (geolocation and status);
- Controlling the exercise by commands to live O/Cs on the range and by remotely controlling devices to produce special effects (such as simulated artillery explosions);
- Continuously collecting and logging data on the state of all players and important events that occur on the range;
- Scoring combat engagements in real time, and kill notification.

A military unit that is to participate in a training exercise on a live range must bring all its combat and maintenance equipment for use in the exercise. This includes all its standard tactical communication gear. The instrumentation and communication equipment needed for monitoring,

control, and data logging associated with the training range is independent of the unit's tactical equipment and is carried by the players *in addition* to their normal equipment.

In order to carry out the above listed functions, a training range, like a test range, requires a number of essential components. The two key differences between test and training range requirements are:

- Instrumentation for accurate scoring of (simulated) combat engagements,
- The absence of telemetry from an SUT.

Monitoring and Control Functions

The monitoring and control of a training exercise is very similar to that of a T&E exercise, and so we use the comm requirements for T&E monitoring and control as representative of those requirements on a training range. That is, a battalion vs battalion training exercise with approximately 2,000 live entities on the range will require an aggregate data rate of about 0.5 Mb/s to be collected at the RCC, including voice circuits and a few imagery links. The RCC gathers the data and sends the control messages, and so a star topology with the RCC at the center provides the requisite connectivity for these functions.

Summary of Comm Needs for a Single Training Range

The monitoring and control functions require digital messaging, voice, and some amount of compressed imagery to be collected by the RCC. For a battalion-sized training exercise, the aggregate communication capacity needed is approximately 0.5 Mb/s, which includes only a few compressed images per minute collected by the RCC. Increasing the amount of imagery collected will result in a corresponding increase in the communication capacity.

The scoring of combat engagements on the training range will require digital messages to be transmitted between players on the range and the RCC with low latency. The peak rates for these messages will depend on the peak firing rates during the exercise. A very conservative estimate of 100 shots per second at the peak of a simulated battle would result in a data flow of less than 150 kb/s in either direction. The required latency over the comm network will be constrained by the requirement to have the scoring appear *instantaneous* to the involved combatants. This limits the latency for message transmission to 300 ms minus the processing time to determine the results of a shot. The accuracy of the sensors which are to provide the data on weapon aiming is a critical issue for a scoring system which must maintain credibility with the players, but this subject is beyond the scope of the present discussion.

As in T&E activities, a critical requirement, and perhaps the most basic, is the need for a communication system that can provide connectivity between the RCC and entities anywhere on the range. The difficulty of obtaining this coverage depends on the size of the geographic area and type of terrain on the range. Direct LOS communication will not be possible in many, if not most, exercises of importance. Thus, even on a single exercise range, the options for complete coverage

of the range from the RCC requires consideration of comm relays of some type. The primary options are UAV relays, satellites, and ground-based relays providing multi-hop comm paths. An operational training range usually has an installed communication infrastructure that includes towers for pickup and relay of messages from anywhere on the range to the RCC. As with T&E ranges, however, the communication challenge is the unprepared range.

2.1.5.4 Summary of Potential T&E and Training Applications of the WI

As in the case of T&E applications, there are two obvious applications for the Warfighter's Internet in training exercises:

- Supporting non-LOS comm between entities on a range and its RCC, and
- Providing connectivity between ranges in a multiple range T&E or training exercise, or connectivity between a T&E or training range and another global network like the DSI.

Regarding the first application, a single comm relay in a UAV could provide non-LOS communication for *all* comm functions within a single test or training range. The services desired include digital messaging, digital voice (telephone type service), and high rate data streams for compressed imagery (and telemetry, if necessary).

The second potential WI application in training exercises or T&E is providing the connectivity between ranges in a multiple range exercise, which could also include a virtual battlespace as one of the nodes. Each node needs to transmit a data stream to each of the other nodes participating in the exercise. A network that could be rapidly deployed anywhere in the world providing coverage over 50 x 50 km areas with a single net entry point would not only be valuable, it could be critical to future joint service training operations as well as T&E at unprepared sites.

2.2 INFORMATION ARCHITECTURE

It is necessary to carefully distinguish between a communication architecture and an information architecture. The former concentrates on what is possible via communications, the latter on what should be provided based on operational policy, command structures, need to know, and implications of security breaches. The goal is for the communication architecture to allow communications between any two points in the global tactical and strategic networks (in addition to commercial networks). However, there are numerous reasons to limit access to specific information. The information services manager will set the policies and the implementation mechanisms to restrict access to only the information that is necessary to execute the tactical mission. As an example, this would include security firewalls and access controls associating legitimate users with specific repositories of data.

Information architecture also deals with the way information is to be accessed in the Warfighter's Internet. Will information be delivered by "push" (automatic delivery) or by "pull" (requests)? Will client-server, collaborative planning, e-mail, etc., applications which are

dominant in the Internet (as used both by civilians and military) be extensible to the Warfighter's Internet? How will voice and video (streaming traffic with strong latency requirements) be delivered over a common transport? How will intermittent connectivity of users (due to battery limitations or due to engagements) affect the delivery of information?

There is one dominant fact concerning the mobile WI. It, like all highly mobile wireless systems, will be capacity-constrained (by available frequency allocations, by payload weight & power in the aircraft, etc.), and therefore information architecture models suitable for wired, wide-bandwidth networks will have to be modified or used with caution in the WI. While the concept of "information push" may be very effective for very wideband channels, it is less attractive for low bandwidth channels where efficient utilization is mandatory. The potential problem with information push is that the user community may be getting much more information than can be absorbed by individuals with different tactical goals. It still makes sense to use information push to alert users of potentially relevant information. This could happen via a mail mechanism. If the mail message is deemed relevant, "information pull" via a client-server or ftp-type service could be used to get the information. File transfer protocol (ftp) would download the entire file. A client-server application would allow the user to browse the document and, if a table of contents is available, only download the relevant information. In the theater, a service like GBS would use a wideband information push strategy. An effective secondary GBS distribution strategy over the mobile WI would require that the information be tailored for relevance and then perhaps use a combination of mail, paging, and information pull.

Although we usually think of the Warfighter as a client downloading information from a remote server in response to small request messages (a data-asymmetric operation), it is also possible for the Warfighter client to transmit threat information and observable damage assessments to the server databases. Other users (planners and Warfighters) then could factor this information into subsequent tactics. While sensor data has not yet been addressed, it should be clear that low rate, periodic sensor data could very well use the same services as are provided to the untethered Warfighter.

The Internet paradigm is a common model for modern open, global data networking. The Internet is built on an IP layer which has allowed interworking across vastly different types of subnets. This will be the basic network level technology used in the WI. However, IP will be changing with the introduction of IPv6, so the Internet will be carrying voice and video with specifiable QoS. The mobile WI implementation will reflect evolution of the Internet, and of IP in particular. The biggest benefit of IP is that it provides a common underpinning for distributed applications, both legacy and future. Legacy applications absolutely have to run (virtually unchanged) and keeping the basic IP network layer allows this to happen. If the communications requirements are modest, then the relevant legacy applications should be able to run with no changes across the wireless, mobile WI. If the legacy application requires "too much" bandwidth then either it cannot be run or it would have to be modified. Modifications could include the use of smart agents that reside in areas of greater communications and computation capability; these agents could, for example, take over the highly interactive signalling between a client and a server, and only report the final information product over the WI to the end user, thus preserving the precious WI bandwidth.

The development of these modified applications or new distributed applications should factor in the reality of limited bandwidth as well as dynamically changing available bandwidth. Applications programming interfaces (APIs) need to be bandwidth-aware and trigger mechanisms that rate-reduce the source information to be transmitted where possible. Alternatively, one can provide special rate-adaptive proxy agents to be placed at the interface between the mobile WI and the higher-rate terrestrial networks to dynamically rate-adapt information passing in the mobile environment.

With the stresses encountered by the untethered warfighters, it is imperative to provide them with as simple and consistent an interface as possible. In today's world, a logical candidate is a realization of a web browser. The goal is to provide as intuitive a graphical user interface (GUI) to the information resources as possible. It is also recognized that today's browsers are largely limited to reading semi-static pages, and that a more generic access and interaction with databases is in the beginning stages. Nevertheless, to state the requirements in terms of a web browser GUI is certainly legitimate; the service to be provided through such a single viewing application is the ability to access many databases of various formats, where the differences are hidden by the "web server" at each of them. The utilization of the web browser is inherently an asymmetric communications load on the WI, since short requests may result in large downloads. This observation is not inconsistent with the actual communications capabilities, since handheld transmitters are certain to be lower in power than airborne transmitters.

3. ARCHITECTURAL CONCEPT AND DESIGN

3.1 INTRODUCTION / OVERVIEW

It has been asserted that the WI should adopt the Internet paradigm of information architecture and should organize the communications connectivity of mobile subscribers (warfighters) in a manner similar to that of cellular telephony networks, with airborne nodes replacing the cellular basestations and wireless backbones between the nodes replacing the wired connections to the Public Switched Telephone Network (PSTN). Internet and cellular-oriented telephony networks are vastly different technologies, but the mobile WI must be built as a blend of the best features of both. The obvious question is why? The answer is fairly straightforward. The rational is that in this manner the Warfighter Mobile Communications Device (MCD) can be made “small” and highly mobile as in cellular telephone systems, while adopting the statistically-multiplexed, connectionless data transfer nature of the Internet for bursty traffic to provide efficient use of the limited wireless bandwidth.

Potentially the WI must support a very large number of subscribers covered by a few airborne nodes. There is no predicting the geographical distribution of users and it is feasible that thousands of users can be covered by the same airborne node. While the majority of users will not be active at the same time, one can be sure that on many occasions many tens, perhaps hundreds, of users will want to participate in active communication sessions. One can build into the airborne node the equivalent of a cellular basestation where each end point in the base station effectively has more communications capability than the individual Warfighter MCD. This asymmetry in capability (and complexity) means that the Warfighter MCD can be made smaller and less complex by allowing the basestation to provide more of the complex processing. Additionally, the basestation provides mechanisms for controlled channel access by the Warfighters desiring communications.

With a point-to-point, circuit-oriented radio architecture, one would need as many dedicated channels on the airborne node as there are active sessions. One would further need a means of matching a subscriber to an available asset on the airborne node. This would require a design modification to the way that point-to-point, circuit-oriented radios typically operate. Even assuming that this modification is executed, it is still clear that the dedicated circuit would normally be severely underutilized for bursty data services. This last point is hard to mitigate without a radically different communications architecture.

A basestation-centric design can provide additional flexibility in terms of channel utilization. A common way to organize any hub-centered (star) network is to assign circuits for the length of a communications session. However, when bursty data is the information to be transferred, this dedicated resource may run at a utilization efficiency of much less than 1%. For voice, the efficiency may be as much as 40% (which is good). However, if data is to be the bulk of the traffic (actually the equivalent in session connect time), then circuit-oriented assignments are a very inefficient choice. This is why statistical multiplexing schemes have been developed. Connectionless systems are very efficient in statistical multiplexing of bursty traffic and the

Internet is the prime example. Today the Internet does not handle streaming traffic (such as voice) well; IPv4 does not have built-in provisions for QoS specification that would serve streaming traffic appropriately. On the other hand, the evolution of IP to IPv6 will provide QoS specifications for several types of traffic, including voice. But whatever protocols are used in the WI, they must provide for both bursty and continuous traffic, and they must provide the service on demand (no dedicated circuits waiting to be filled).

Connection-oriented wireless telephony base stations (like commercial cellular telephone systems) include a circuit switch to set up an end-to-end route for the traffic. The Internet, a connectionless wired system, reduces all communications to a packet-oriented basis and utilizes a series of routers to direct each packet along its own way to the destination according to the IPv4 service model. The Warfighter's Internet should not be equated to the connectionless operation of the current Internet IP service model. IP evolution to IPv6 will allow sessions to ask for a requested quality of service (QoS) for both bursty and continuous traffic which is needed for the Warfighter's Internet, but the Warfighter's Internet has additional protocol needs not included in IPv6 that must handle the routing problems due to the mobility of the airborne nodes and subscribers, as well as multicast and QoS performance in the face of the unreliability of the wireless links due to propagation variability.

The connectionless mode of operation is a way to make channel utilization efficient for bursty traffic. The following sections look at additional ways to provide efficient operations depending on the particular channel being considered.

The strawman architectural concept describes how each subsystem implements mobility, QoS, multicast, efficient channel utilization, user access with priorities, etc.

3.2 INTERFACES, BOUNDARY OF WARFIGHTER'S INTERNET

The Warfighter's Internet is a network that must interface with other networks and systems found in the tactical theater and be able to pass information to and from them. In fact, the Warfighter's Internet must provide global reachback to CONUS information systems to allow transfer of information worldwide. Within the Warfighter's Internet, special transmission protocols can be used, but at the boundaries of the Warfighter's Internet, protocols appropriate to the external network(s) must be used. There are several possible entry points to the Warfighter's Internet. They include at least one major entry node (or gateway) which provides connectivity to the theater comms (MSE, etc.) and the DISN; through this point, worldwide connectivity and theater connectivity is provided (see Figure 3-1). This entry point is on the Warfighter's Internet backbone, and will be on the end of a high data-rate link from an airborne node, either directly from the aircraft or over a satellite link (shown in green to indicate it is not part of the WI backbone, but transparently connects an airborne node to the DISN STEP, which could be an alternative entry node). A second type of entry point is at an individual subscriber to the Warfighter's Internet; a variation on this type is where the entry point is a traffic concentrator connected to a local area network (LAN) (possibly a wireless LAN serving a number of individual users). Such concentrator points can also be a large terrestrial command center, a Navy command ship, or possibly an Army Radio Access Point (RAPs). (Such concentrator points could actually

be reached by direct extension of the WI backbone if capacity needs require it.) The final entry point type is directly on the airborne node; it is present when there are other radios on the aircraft that connect users to the Warfighter's Internet. This entry point will be found on the Airborne Communications Node (ACN) which carries a substantial number of legacy radios, along with gateways to convert their signalling methods to that of the Warfighter's Internet format.

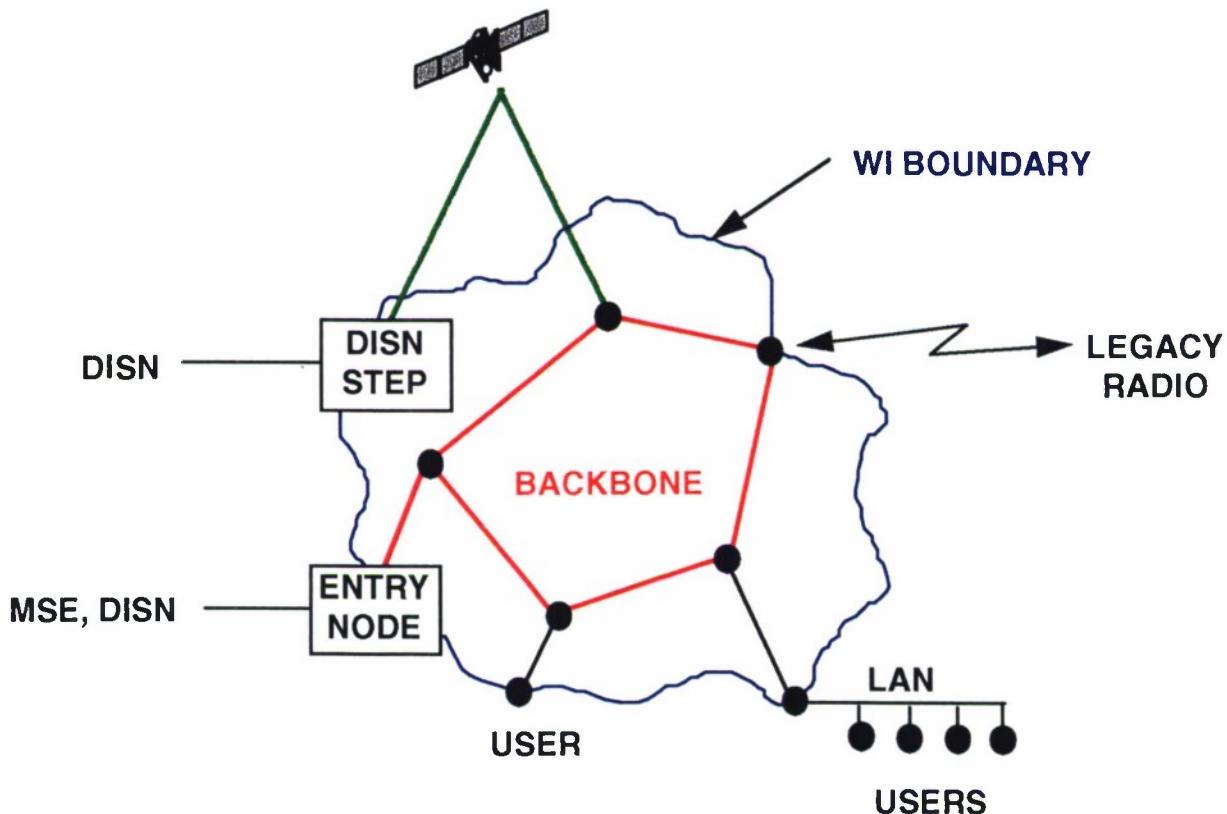


Figure 3-1. Interfaces and boundary of the Warfighter's Internet.

A particular issue is in the form of the “packing” of information presented to the Warfighter's Internet boundary. One standard format is IP packets; the packet header contains (among other things) an IP address identifying the destination of the packet. Another format which is of interest is ATM (Asynchronous Transfer Mode); here the transmission quanta is the ATM “cell” (53 bytes) instead of the packet, and the routing of each cell is determined by its header, which identifies it as belonging to a particular end-to-end virtual circuit. ATM traffic may include IP sent over ATM signalling as well as pure ATM (for carrying streaming traffic). ATM formats will be present at major nodes of the MSE (see Figure 3-2), including Node Centers and some concentrator nodes such as the RAP (Radio Access Point) equipped with the HCTR (High Capacity Trunk Radio) in which IP traffic on wireless subscriber networks (e.g., Tactical Internet) is converted to/from ATM format. In other cases, end users present (or expect) IP packets to (or from) the WI boundary. If the Warfighter's Internet is to provide interconnectivity of IP and ATM signalling presented to it, then there must be a place for conversion between ATM and IP.

This fact is independent of what signalling format is utilized internal to the WI (which has not been discussed yet). A general solution is to convert to/from the internal signalling format at the boundary of the WI (at a terminal or on the aircraft, as identified in Fig. 3-1). Of course, if either ATM or IP is chosen to be the internal signalling format, then there will be cases where no conversion is needed at a boundary, but other cases where the conversion is necessary. But just carrying the information through the WI successfully is insufficient; the end-to-end protocols (of whatever type), including signalling protocols, must also function properly. While today IP is associated with connectionless packet service and ATM is associated with streaming cell traffic, the evolution of IP and ATM will eventually allow either type of traffic to be carried under control of either signalling method.

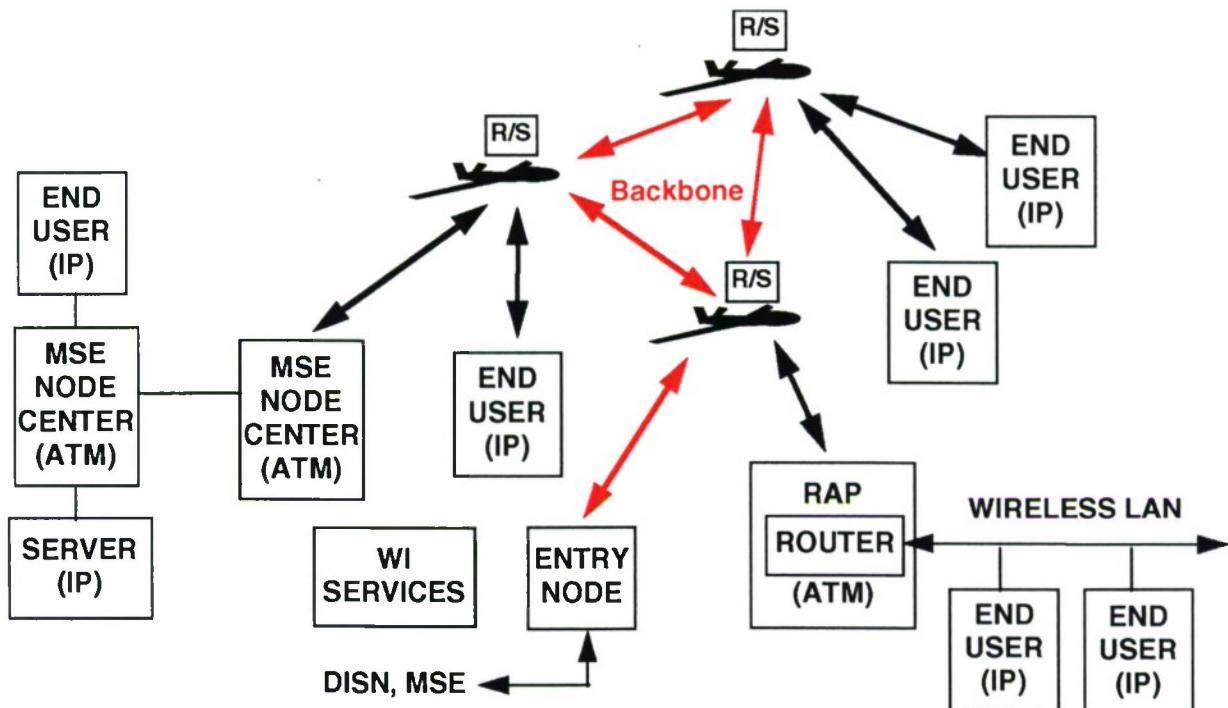


Figure 3-2. Warfighter's Internet network architecture.

3.3 EFFICIENT COMMUNICATIONS

Wireless systems present severe challenges in comparison to wired systems beyond just limited bandwidth, and with the widespread introduction of high data rate fiber cables the relative bandwidth disadvantages are increasing. The main compensating advantage is user mobility, and this is clearly a huge operational advantage in many tactical environments. However, the desires of the users in terms of applications is being conditioned by what they see on workstations on a wired network while in garrison, that is, in communications-rich and computationally-rich environments. The challenge of the mobile WI is to approximate the utility of some of these applications in deployments with communications-poor and computationally-poor environments.

For a number of reasons, existing civilian mobile communications systems are organized in a cellular fashion, and cells are tied together with a backbone. However, even the commercial mobile cellular wireless environment faces some severe physical limitations. For both safety and portability reasons, the user handsets are restricted to low transmit power levels. Since the user handsets have omnidirectional antennas, multipath degradation becomes a limiting factor. In cellular systems accommodating large numbers of users, frequency allocation is also a limiting factor, and elaborate schemes have been designed to allow reuse of bandwidth so that the aggregate capacity is reasonable. In spite of these advanced designs, communications bandwidth is the most difficult and expensive asset to acquire.

When a cellular communication organization is applied to military situations, there are further constraints on bandwidth. Jamming reduces communications forcing tradeoffs between available data rate throughput (capacity) for jamming protection. Moreover, there is some additional flexibility in commercial cellular systems: one can subdivide cells (geographically or spatially) and reuse the bandwidth. This strategy is not always practical with airborne platforms: more aircraft may not be readily available and those that are may be forced to service large footprints.

The fundamental fact remains that it will be necessary to live with limited communications capacity, but the available resources can be used to maximize the throughput seen by users by applying two design rules:

1. Utilize the available capacity in the most efficient manner possible.
2. Write applications that require the minimum amount of communications capacity.

These are the two considerations that will drive the communications architecture and the information architectures, respectively. Figure 3-3 shows an overall approach on how to attack both limitations and still deliver the type of services that will make an operational difference to the Warfighter.

The next section is devoted to providing an overview of the network with particular attention to mechanisms for better utilization of all of the links. While it is difficult to increase raw channel capacity, it is possible and mandatory to seek out every opportunity to improve available bandwidth utilization, which will be perceived by the user as improved throughput.

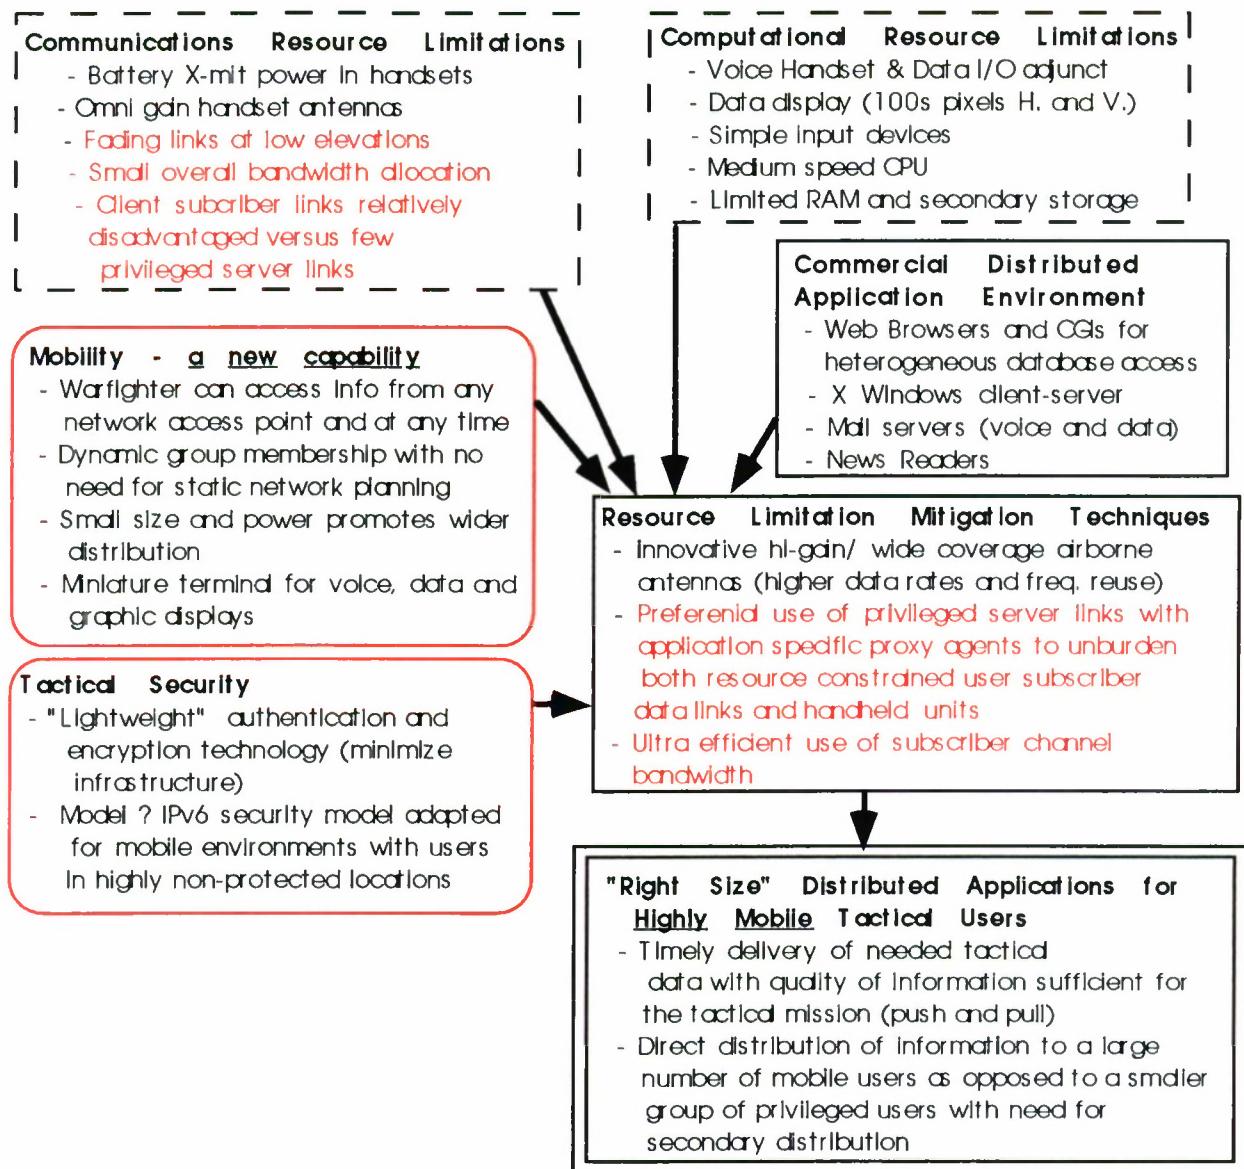


Figure 3-3. Providing high impact tactical distributed applications in a resource-constrained environment.

3.4 COMMUNICATION ELEMENT DESIGN CONSIDERATIONS

There are three main communication elements which need to be considered: uplinks and downlinks to individual subscribers, e.g., a dismounted soldier; uplinks and downlinks to more capable terrestrial nodes, e.g., ships and Army RAPS; and the airborne-airborne cross-links which form the WI backbone, including the link from an airborne node to the terrestrial entry node. Each of these will be discussed in turn below.

3.4.1 Uplinks and Downlinks to Individual Subscribers

3.4.1.1 Introduction

The communications capacity of the uplinks and downlinks between the airborne platforms and the individual subscribers are the most limited resources of the Warfighter's Internet. Not only must these resources be assigned carefully on the basis of need, but the signalling techniques used on these wireless links must make the most efficient use of the available data rate and frequency allocations. This design optimization must take into account the nature of the communications traffic: it will be predominantly data, with some voice and other constant rate traffic (e.g., video). The data traffic will be bursty computer communications. Computer traffic is characterized by bursts of activity interspersed with relatively long periods of silence, so that the channel is actually in use by a given user only a small part of the time (perhaps < 0.1%). It would be very inefficient to assign a permanent, dedicated RF channel (a circuit) for this purpose. Another consideration to be taken into account is the use of the Mobile Communications Device (MCD) as a terrestrial concentrator for a wired or wireless LAN. When used in this capacity, the uplink channel may at times contain the aggregate traffic from multiple users.

It is useful to examine the downlink and uplinks separately when considering the efficiency of channel utilization as shown in Figure 3-4. The downlink is the simpler case. The downlink will consist of a stream of interleaved packets or cells directed to many different users sharing this downlink beam. Since the communications will have both different real time demands and user priorities, the UAV's downlink traffic manager can efficiently schedule and interleave individual packets transmissions so as to match the time demands and the priorities. The traffic manager can completely fill the downlink so that lower priority excess traffic will be delayed until the higher priority traffic is serviced. A model for the downlink from the airborne node would be that of a (wireless) LAN, where all of the traffic for all of the users under the aircraft coverage would be broadcast sequentially in packets. End users would listen to the downlink stream and pick off the packets intended for them. Because of the low utilization rate of any individual user, many users normally share this downlink LAN and receive responsive service. Since the downlink is a broadcast, a further efficiency improvement can be realized when the same message is being transmitted to multiple users. If a multicast (or broadcast) address is used, the message need only be transmitted once and all the active intended recipients can collect the message at the same time.

Using the uplinks efficiently is a much more challenging problem. This is a result of the uplink users being relatively numerous with respect to the available receiving resources in the airborne node and thus efficient contention for and assignment of these communications assets is the prime problem. End users must share the available uplink resources by only using them when there is actual data to send. This implies that there will be some variation of demand access to the uplink channels. The difficult problem is to be able to predict reliably the actual data traffic profile. One solution is to reserve a "circuit" for the duration of a session. As noted previously, this is a very inefficient process, particularly for data communications. The other extreme is for the different users to transmit uplink packets (or cells) as soon as they are generated without scheduling exclusive use of an uplink asset. However, this method can result in different user's packets (or cells) colliding and thus not being received. The objective of any efficient uplink

access scheme is to maximize utilization of the available RF resources, while maintaining the desired Quality of Service (QoS) for each user. The technique used could vary from pure ALOHA (transmit a packet/cell anytime when you have one) to variations on slotted and/or reservation techniques. In following sections, example implementations are discussed for this uplink assignment problem.

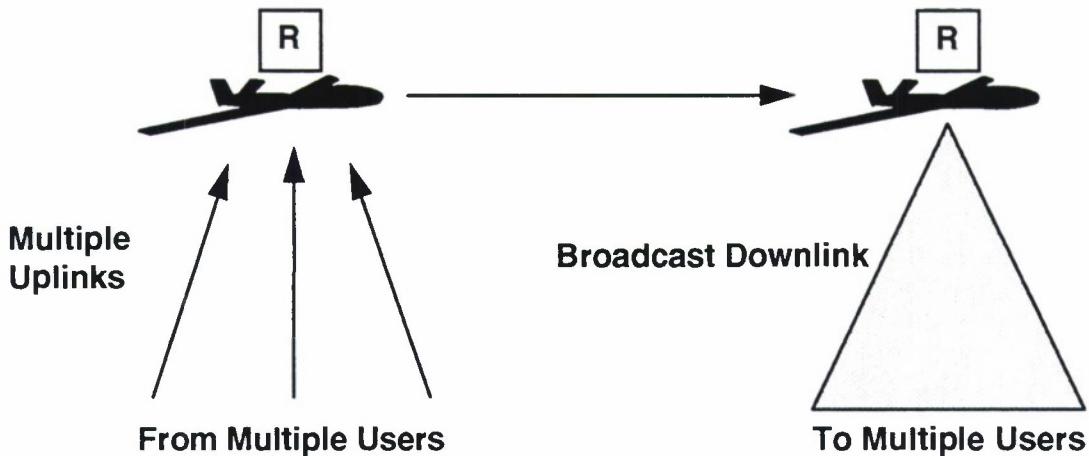


Figure 3-4. Efficient use of uplink and downlink resources.

3.4.1.2 Design Criteria

To be most efficient in sharing resources it is advantageous to merge requests for capacity to draw from as large a pool of resources as practical. This will be more efficient than, say, dividing users and resources into partitioned subsets, e.g., channel by channel. (However, it may be desirable to divide services and resources into broad categories if transmission needs are highly heterogeneous, e.g., bursty datagram and streaming voice.)

The gain in efficiency by pooling resources is illustrated in Figure 3-5. The figure shows the probability of message blockage as a function of the loading on each of a number of communication channels (uplink demodulators or downlink time slots). The number of channels pooled is the parameter for each curve. It is seen that the blockage probability decreases dramatically as the number of channels in the pool increases even when the offered load *per channel* is fixed. From this it can be seen that treating uplink and downlink resources as a shared pool to draw on dynamically in response to demand is advantageous. (A simple M/M/m traffic model was used only for illustration with blocked traffic held in the system without the need for repeats. This model assumes that messages arrive randomly and independently of each other. The message lengths have an exponential distribution. Also, the number of users in this model in effect implies that there are many more users than channels.)

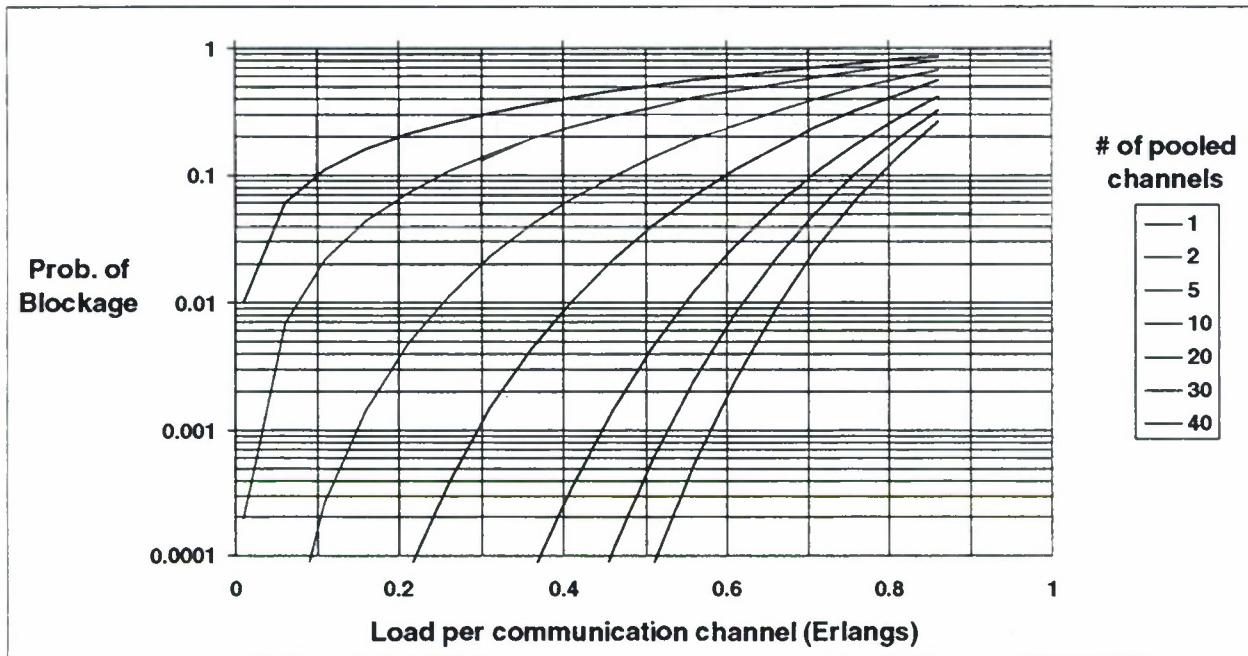


Figure 3-5. Illustrating the advantage of pooling communication resources (M/M/m model).

In a number of radio systems, the incoming and outgoing channels are similar. This is the general situation with packet radios. However, a cellular-oriented system like the mobile WI (or any star or hub-oriented radio system) is considerably different, and for the WI there are two main differences. First, the average transmitted user power is limited to about 3 Watts while the transmit power on the UAV can be orders of magnitude higher. Second, it is extremely difficult to manage the efficient utilization of the aggregation of uplinks. These two facts are the basis of the assertion that the uplinks are the basic system level communications constraint, and if means can be found to manage the uplinks efficiently, then the overall system will be efficient since the downlink and cross-links can easily be managed efficiently. In fact, to a high degree the overall Quality of Service (QoS) of delivered communications will depend essentially on how the uplinks are handled.

The communications data rate that can be supported on both the uplinks and the downlinks depends on transmitted power, frequency choice, and antenna type. The desire is to operate to forces on the move, which implies simple omnidirectional antennas on small receivers on the ground and therefore operation at frequencies below 2 GHz, as determined by allocations. The antenna on the aircraft can have a pattern shaped in elevation that puts most of the transmitted power (or receive gain) toward the horizon (but is still uniform in azimuth) to make up for the substantial difference in distance between users immediately under the aircraft and those at a distance of, say, 100 miles. Under this condition, a 50-Watt transmitter in the aircraft can deliver a T1 data rate (1.544 Mbps) to the small receiver. Likewise, a 3-Watt ground transmitter can support 64 kbps on the uplink. Both of these calculations (found in the following section, 3.4.1.3) include 15 dB of margin allocated to propagation anomalies such as fading.

Note that these data rates are burst rates; because of the burstiness of computer data, it is not expected that any single user would operate at anything near these rates on the average. These characteristics are illustrated in Figure 3-6.

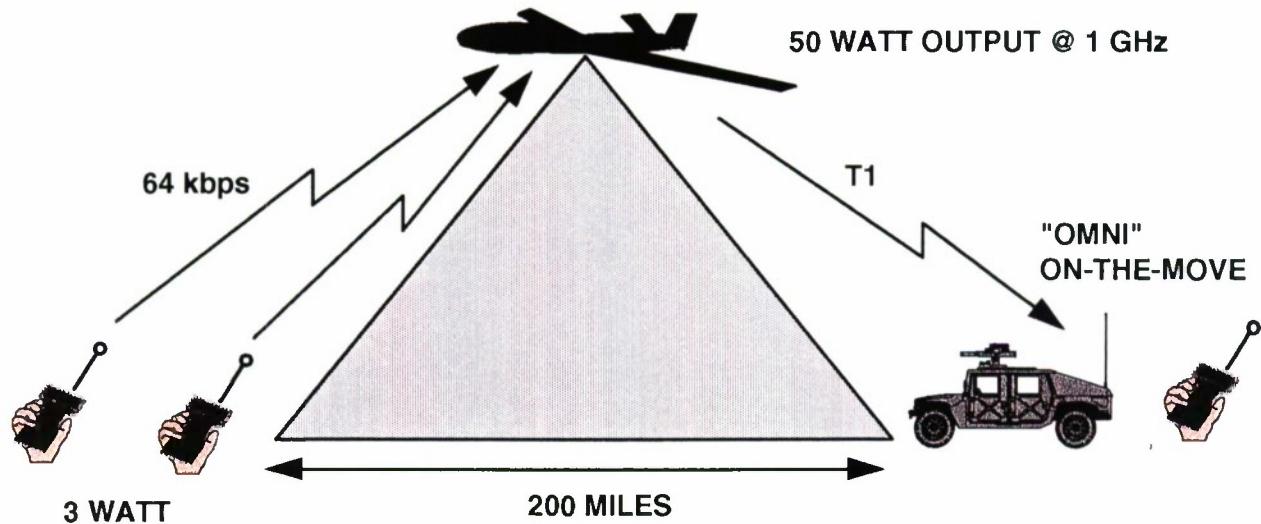


Figure 3-6. Uplink and downlink data rates to and from users with omnidirectional antennas.

Having established the basic principles of the uplinks and downlinks, the next question is that of resource control and allocation. It is assumed that users on the ground make contact with an airborne node over a signalling “channel” (which may be a separate physical channel or embedded in the multiplexed data traffic); when authorization has been established, the aircraft subscriber controller responds with an uplink channel assignment sent in downlink time slots reserved for the control function. The controller also learns the user’s identity (IP address or system ID) and passes the association of the user with this aircraft to a Mobility Data Base for future routing use. At this point the user is able to exchange communications traffic with the aircraft. The routing functions in the aircraft determine whether the recipient is reachable from this aircraft or if it will be necessary to route the traffic over the WI Backbone to reach the recipient. Likewise, the routing function monitors traffic arriving on the Backbone that is intended for any user registered with this aircraft; any such traffic is then transmitted on the downlink from the aircraft. In a similar way, any traffic that arrives through another radio on the airborne platform can be introduced into the WI system by treating it as if it had arrived over the cross-link. This model is particularly applicable to the ACN which has many legacy radios on board. Of course the same model applies to traffic to be sent out through a separate radio on the airborne platform.

Although in the discussion so far, it has been assumed that the subscribers are homogeneous (in antenna gain, transmitter power, link conditions, and communications requirements), this need not be the case. Larger ground terminals (concentrator nodes) could have more antenna gain or transmitter power to increase their data rate significantly. Alternatively, uplink transmission rates could be increased by the use of multiple, narrow beam

receive antennas on the aircraft (which are also useful for rejection of unwanted jamming signals). Downlink transmission capacity from the aircraft could also be increased by utilizing multiple (higher gain) transmit antenna beams. In addition, an adaptive physical layer would be able to adjust the uplink or downlink for best performance under whatever link conditions prevail for each user.

3.4.1.3 Uplink and Downlink Performance and Discussion

The Warfighter's Internet will rely on radio link technology being applied or developed by the Airborne Communication Node (ACN) and other programs. Previous ACN studies developed some of the basic link relationships that give an estimate of performance and these developments are continuing. This section highlights some of these results.

The following two plots (Figures 3-7 and 3-8) show the relationship between data rate, design range and frequency selection. Examples using a 50-Watt airborne node transmitter and a 3-Watt subscriber transmitter are given. A nominal 15 dB of link margin over free space is

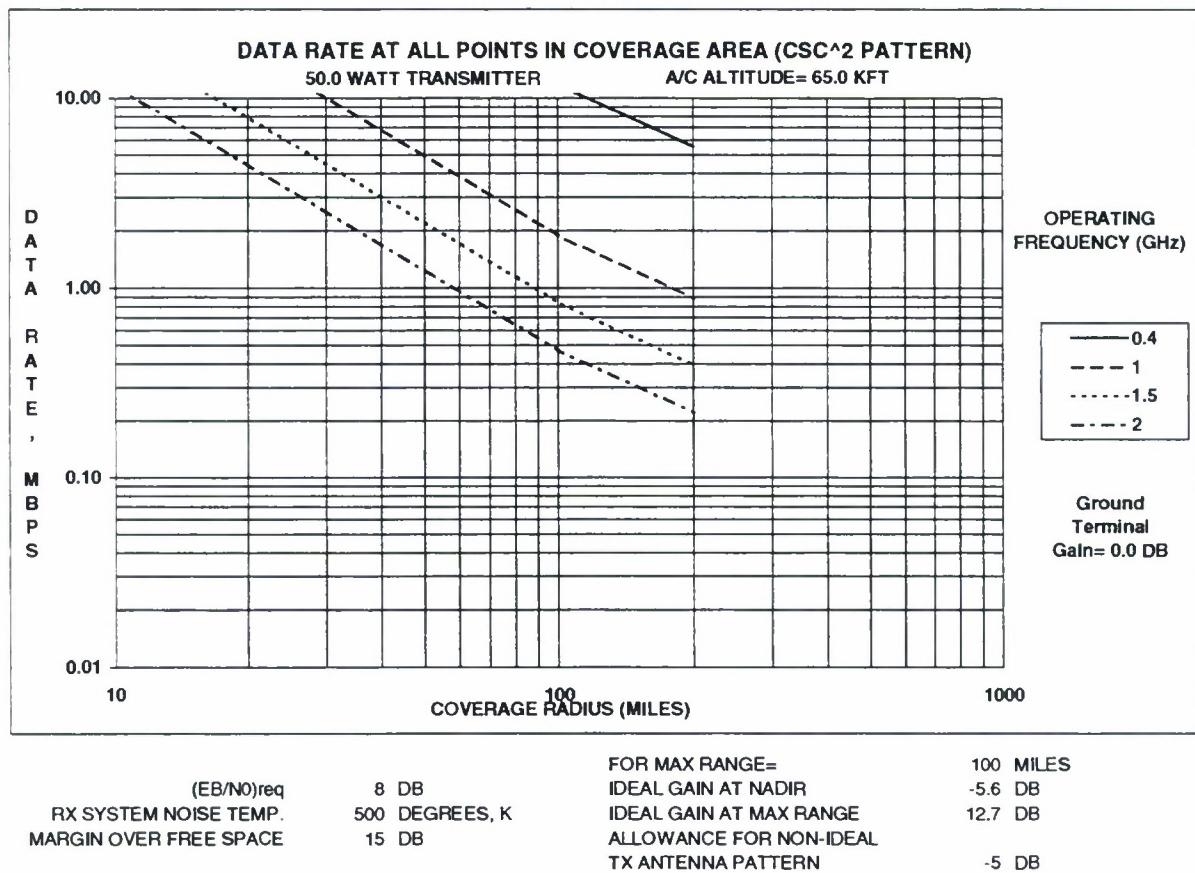
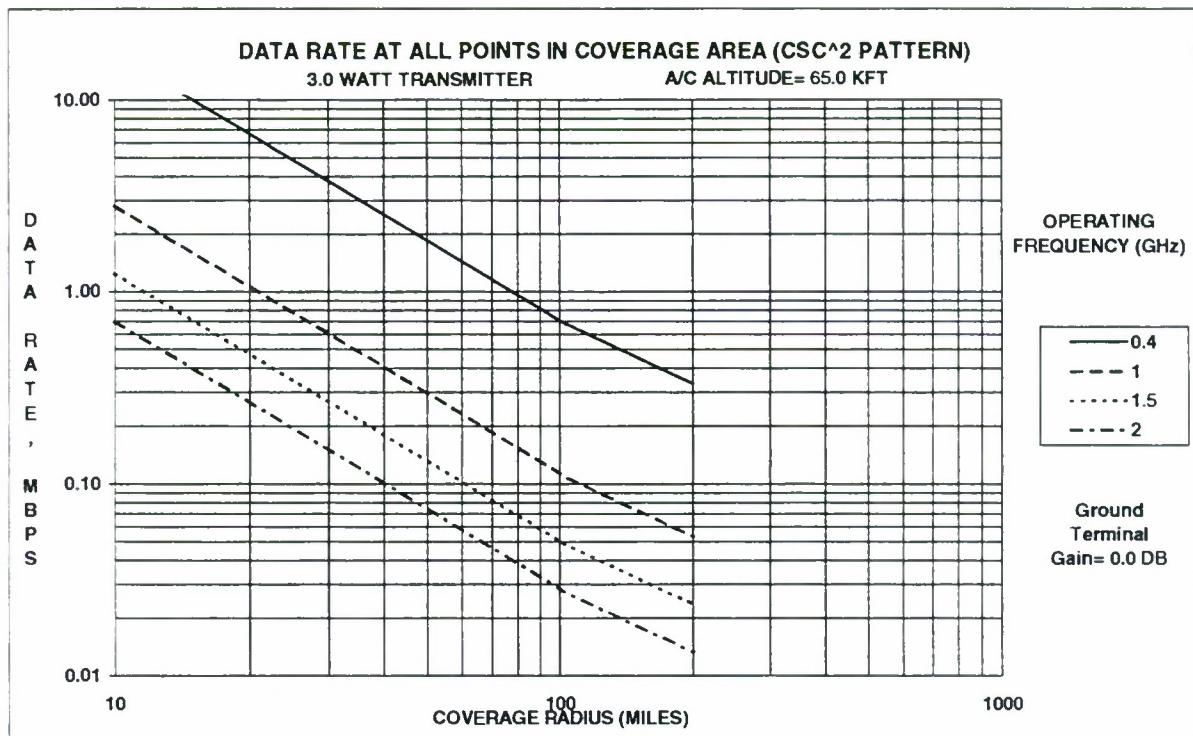


Figure 3-7. Nominal airborne node-to-individual subscriber downlink performance.



(EB/N0)req	8 DB	FOR MAX RANGE=	100 MILES
RX SYSTEM NOISE TEMP.	500 DEGREES, K	IDEAL GAIN AT NADIR	-5.6 DB
MARGIN OVER FREE SPACE	15 DB	IDEAL GAIN AT MAX RANGE	12.7 DB
		ALLOWANCE FOR NON-IDEAL TX ANTENNA PATTERN	-5 DB

Figure 3-8. Nominal individual subscriber to airborne node uplink performance.

included to account for blockage, attenuation, and multipath effects. Only background noise is taken into account; multiple access and other forms of interference is assumed to be negligible. (The validity of the assumptions about the benign nature of other forms of interference will need to be revisited when specific radio system designs are evolving.) It is important to add that individual users will experience wide variations in link conditions and it would be highly advantageous for the Warfighter's Internet to include adaptive features which would permit optimum use of whatever signal-to-noise ratio is available. The plots assume a csc^2 airborne node antenna pattern which compensates for the wide variation in user range from directly below to about 100 miles away. (Note that the data rate given is available throughout the indicated design range, with a different antenna design implicit for each range.) In practice it will be difficult to achieve such a pattern and so 5 dB has been subtracted from the ideal. As an example, the airborne antenna designed for a 100 mile range is assumed to have 7.7 dB gain at its furthest range and -10.6 dB at nadir. The need for this type of pattern, which requires a vertical aperture of about 4 wavelengths, is less critical at lower frequencies where the free space path loss is less.

While it is desirable to keep the airborne node antenna simple, in an operational system there may be the need to provide more antenna gain for: greater frequency re-use for more

capacity, protection from interference and jamming, and more link margin to counter propagation effects. This may lead to the need for sectored coverage areas with a multiple beam antenna (akin to those used in the cellular industry). The Warfighter's Internet routing algorithms must be extendible to accommodate this eventuality. An important aspect of this will be the need to provide for handover from sector-to-sector.

The nominal 15 dB link margin will be particularly problematic for users at the furthest ranges and correspondingly lowest elevation angles of propagation (see Figure 3-9). The plot below also shows the unfortunate geometric fact that most users are likely to lie at low angles assuming a reasonably uniform distribution of terminals. Because of the difficulty in operating at low elevation angles, until some experimental experience has been obtained, quantitative predictions of grade of service will be difficult to estimate accurately. Some preliminary estimates can be obtained from PCS/cellular and MSS (mobile satellite system) industries. Foliage may be a particular concern to the Army. A simple model of foliage attenuation (Weissberger) that is being considered as an industry reference illustrates the problem. For the conditions shown in Figure 3-10 (10 meter stand-off distance from 10 meter tall trees), it is seen that foliage alone can cause considerable attenuation at low elevation angles, although it is less of a problem at lower frequencies.

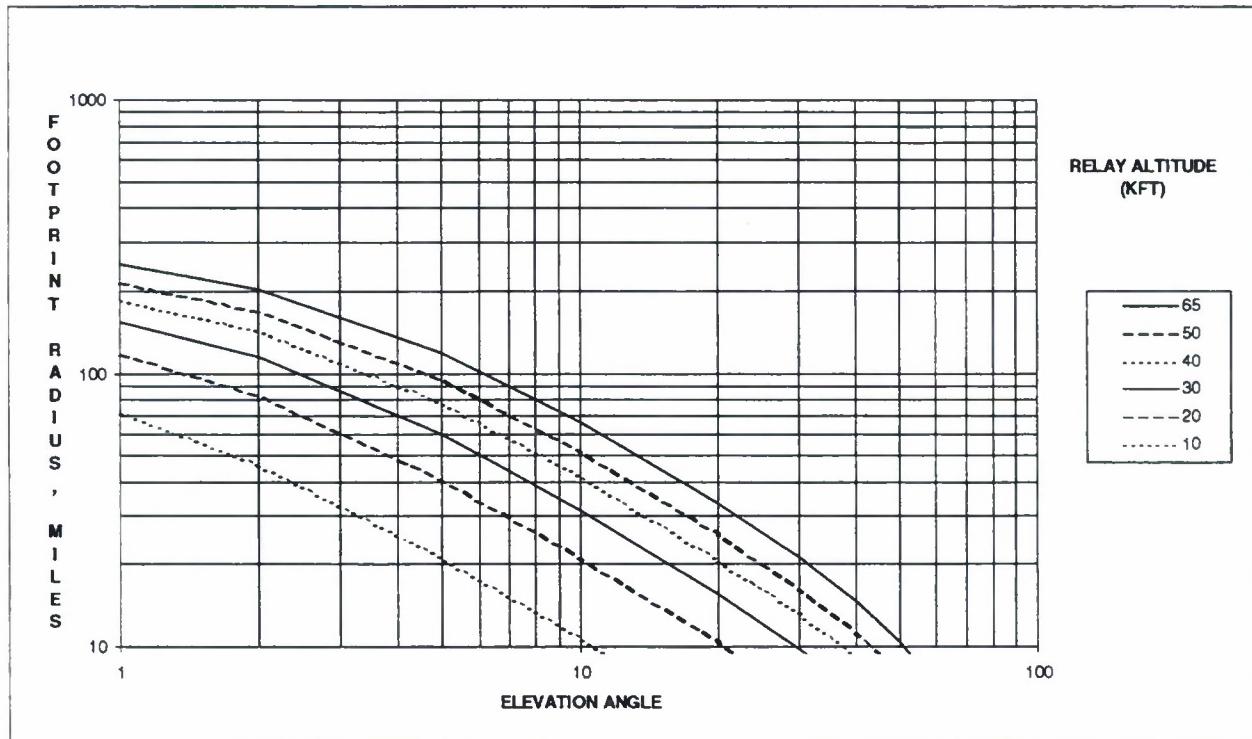


Figure 3-9. Footprint coverage radius (n miles) vs elevation angle (degrees).

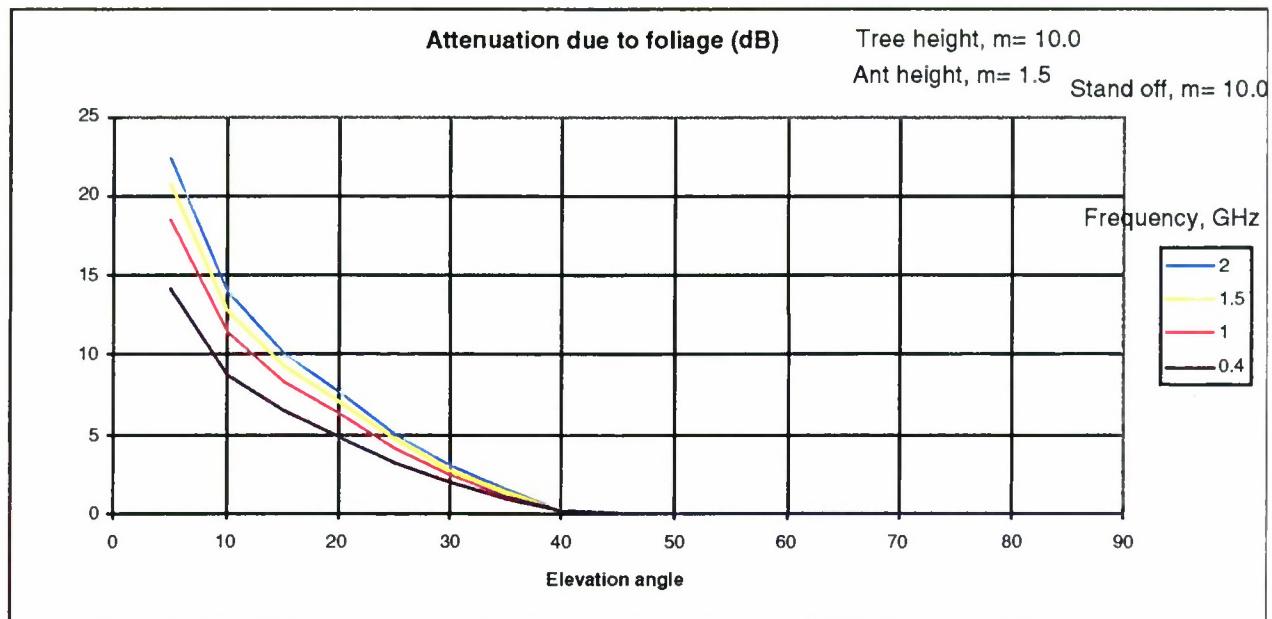


Figure 3-10. One foliage scenario.

The selection of operating frequency will be a critical parameter in determining link performance. Most factors point to the use of a frequency in the low UHF band (around 300 MHz) as best although it is reasonable to consider frequencies up to about 2 GHz. This can be noted from the data rate vs. range plots as well as the foliage attenuation model. However, it can not be assumed that spectrum will be available at this (or any other) frequency of interest—an issue that is being actively addressed by the ACN community. One factor that argues away from the lower frequencies is the difficulty of designing efficient, but small, antennas in this regime.

Another factor in determining system performance (particularly capacity) will be the ability to take advantage of frequency re-use, i.e., reusing the same frequency band in different geographic regions by multiple ground users and air nodes. Current analog cellular systems (AMPS) allow frequency reuse by dividing the total band into seven (approximately) equal portions then permitting non-neighboring cells to re-use the same portions. In addition, current systems use sectored antennas at the base stations to reduce the effects of interference. The newer spread-spectrum technologies also use sectored antennas and claim an increase in capacity by permitting neighboring cells to use the same frequency bands, but mitigate interference problems with the spread spectrum. The frequency re-use problem exists on both the uplink and downlink. It may be possible in the Warfighter's Internet to be even more efficient than current practices by taking advantage of the knowable precise position of each user and tailoring frequency reuse assignments accordingly.

- Modulation, coding and spread spectrum

The data rate performance curves presented above assumed that a reasonably efficient modulation and coding scheme will be used, requiring a signal-energy-to-noise ratio, E_b/N_0 equal

to 8 dB. This figure should be achievable by adapting available hardware (chip sets) unless the multipath environment is unusually severe. In any case, it is presumed that the modulation and coding scheme will also include spread spectrum for protection against frequency-selective interference and channel propagation effects. Spread spectrum may also be used as a way to create code division multiple access channels which enable frequency re-use from platform-to-platform or antenna sector-to-antenna sector. Of the two main categories of spread spectrum, frequency-hopping may be preferred in the Warfighter's Internet environment over direct sequence spreading for at least two reasons: its ability to avoid crowded portions of the frequency spectrum and its relative ease in creating orthogonal uplink multiple access channels with synchronized hopping. The anti-jamming performance of both techniques are comparable when spread over the same bandwidth (although direct sequence spreading may be more advantageous by a few dB because it can more easily take advantage of coherent modulation techniques, assuming that a reasonably uninterrupted band of frequencies can be found) and nominal performance is shown below (see Figure 3-11). These results show that it will be difficult to provide large amounts of jamming protection unless significant antenna pattern discrimination is employed by the airborne node. (For example, a jammer-to-signal ratio of 20 dB is tolerable when communicating at 10 kbps and spread over 10 MHz. This leaves a 3-Watt user vulnerable to a 300-Watt jammer if nulling or other antenna discrimination is not employed at the airborne node. The situation would be even worse if the user signal was attenuated and the jammer's was not.)

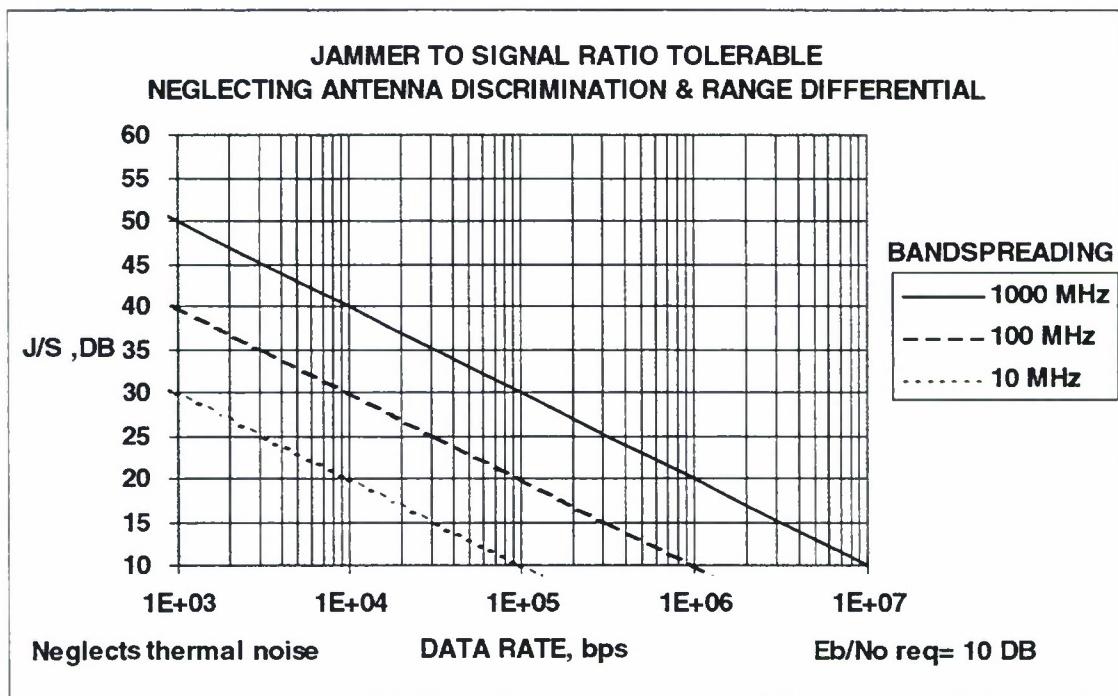


Figure 3-11. Nominal jamming protection.

3.4.2 Uplinks and Downlinks to Terrestrial Nodes

The terrestrial nodes that will be served by the Warfighter's Internet can include, for example, Army Radio Access Points (RAPs), Navy ships, and Air Force Combat Air Operations Centers (CAOCs). It is assumed that these nodes, acting as points of traffic aggregation, will have the requirement for communications with the airborne node at rates to several Mbps on a nearly continuous basis. Their communication characteristics would be more akin to trunked service, although with data rates that can adapt to real-time demand. The communication terminals can be expected to employ directive antennas unlike those of the individual subscribers.

Figure 3-12 shows the results of link calculations performed for the case of an airborne node with a single beam csc^2 antenna pattern (as for the individual subscribers as discussed above) communicating with a terrestrial node having a dish antenna with the diameter shown as a parameter. It can be seen that a modest amount of RF power (25 Watts) operating with a small aperture ground terminal (about 1 foot) can support significant data rates (10 Mbps) to ranges of about 100 miles. To first order, the calculation is independent of operating frequency because the airborne antenna gain is being held fixed as well as the ground terminal aperture. However,

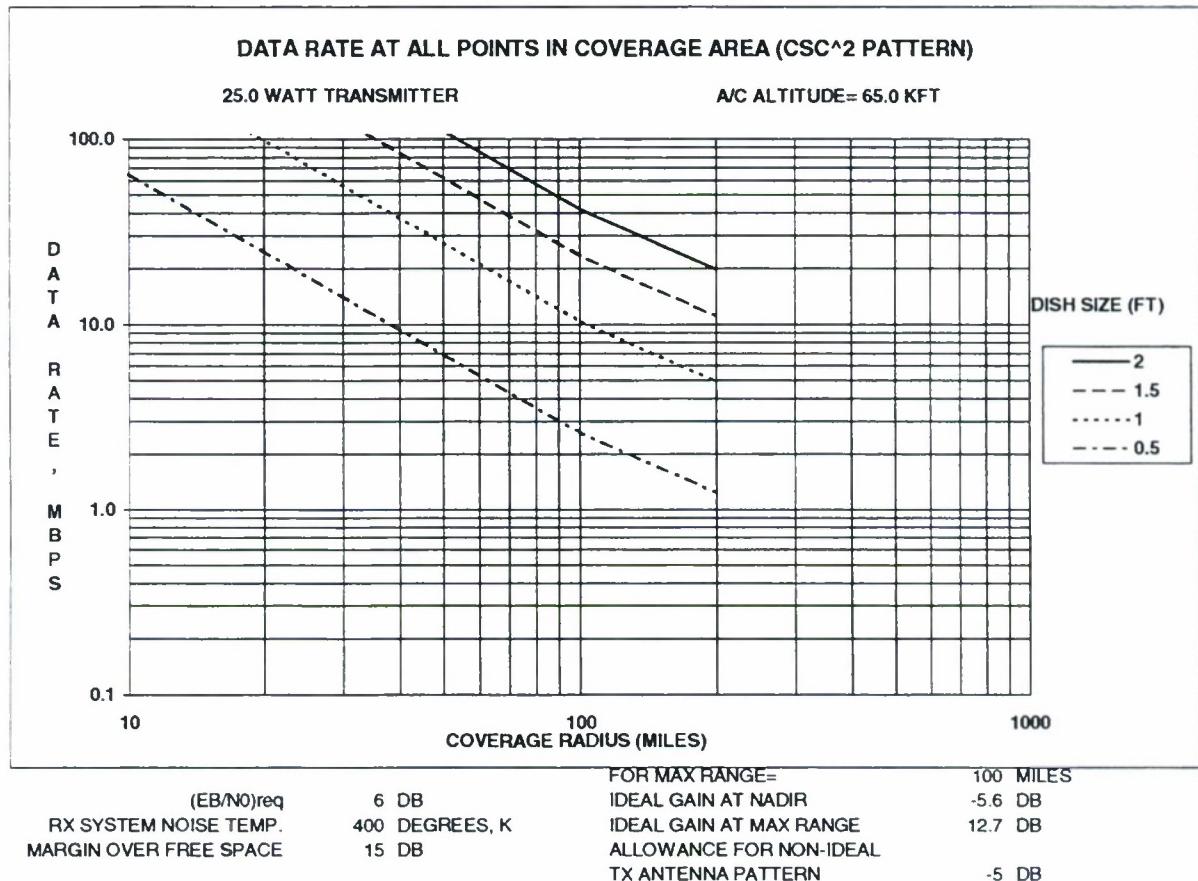


Figure 3-12. Terrestrial node links to/from airborne node.

because microwave frequencies will almost certainly be used in order to obtain a generous amount of bandwidth, the effect of attenuation due to atmospheric gases and rain must be taken into account. Below about 10 GHz these effects are not overly severe. Above 10 GHz, rain is a significant factor and there are specific frequency bands that are subject to large attenuation due to atmospheric gases (water vapor around 22 GHz and oxygen around 60 GHz). (Foliage and other obstructions are also factors, but their effects are uniformly severe across the frequencies of interest.) A desirable band of operation would be in the vicinity of 8-10 GHz, known as X-band (currently used for Milsatcom and the CDL air-ground link).

Operating at higher frequencies, however, does have some desirable properties. In particular, the size of the antenna needed on the airborne platform for a given amount of gain decrease with increasing frequency. This makes it easier to form either specially shaped airborne patterns, including patterns with nulls in the direction of jammers, or individual beams that can form narrow beams on individual ground terminals. The latter would permit intense frequency reuse and can provide extra gain to overcome path losses. On the other hand, operating at higher frequencies also implies the need for more accurate spatial angle tracking.

3.4.3 Airborne Node to Airborne Node Cross-links

The Warfighter's Internet will require a high data-rate, multi-node backbone network connecting the airborne nodes and the WI Entry Node(s). Because the airborne nodes are constantly in motion, the topology of connectivity within the Backbone can frequently change. Such change can be necessitated by aircraft entering or leaving the backbone, by relative motions between the aircraft as they pursue other missions, and by changing coverage needs to support force mobility. This change in connectivity results in routing changes to reach end users served by a given airborne node (or when their service point is changed to another airborne node). The establishment, maintenance, and change of the backbone connectivity and subsequent routing information is a major design issue for the Warfighter's Internet. The changes in the Backbone must be made without human intervention; a self-organizing algorithm must be implemented to match both the link capabilities (allowable data rate, frequency assignments, available antennas, etc.) and the time constant of connectivity change.

Cross-links providing wide area connectivity among airborne nodes are a necessary component of the Warfighter's Internet. In this section we discuss the principles behind implementing these cross-links and some related topological considerations. This discussion will lead to the conclusion that establishing individual air-air links can be achieved with relatively straightforward RF techniques. The more difficult problem will be that of managing this connectivity in the face of the mobility and maneuvering of the nodes.

Engineering the airborne-node constellation is highly dependent on the antenna patterns used. Figure 3-13 shows three basic types of antennas that can be considered. The corresponding patterns range from omnidirectional (at least in azimuth) through multiple narrow steered beams to multi-beam arrays. The narrower beams should be thought of as forming the main data carrying paths. The omnidirectional antennas can be used to provide "discovery" beams and to maintain connectivity during maneuvers or other link degradations, albeit at relatively low rates.

- **OMNIDIRECTIONAL ANTENNA**
 - No pointing required
 - No frequency reuse
 - Lowest data rate
- **MULTIPLE ANTENNAS**
 - Separate steered antennas
 - Must know where to point
 - Can reuse frequencies
 - High data rate
- **MULTI-BEAM ANTENNA**
 - Full coverage from antenna array
 - High gain in all directions
 - High data rate
 - Can reuse frequencies
 - Most complex; multiple Rx, Tx
 - Growth option

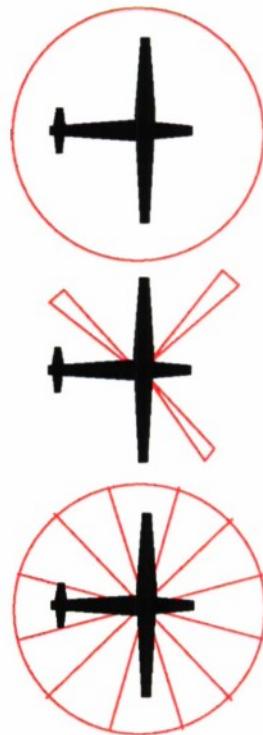


Figure 3-13. Cross-link antenna patterns.

3.4.3.1 Directive Beams

In order to illustrate the feasibility of relatively high rate cross-links, first consider the capability of a cross-link formed with small aperture antennas at each end. An antenna with an aperture of diameter D feet operating at a frequency of f_{GHz} in GHz will form a beam of width $70/(D f_{\text{GHz}})$ degrees. For example, a 9-inch dish used at 10 GHz (as used in the ABIT and CDLS systems) will have a beamwidth of about 9 degrees. This narrow beam will have a corresponding gain (or “focusing” power) of 25 dB. A representative link calculation using these parameters with the assumption of a 25-Watt transmitter is shown in Figure 3-14 along with a data rate versus range plot. (A nominal link attenuation has been included to account for the presence of atmospheric gases. This value is a function of operating frequency. Also, the effects of rain along the path, which can be significant at frequencies above about 10 GHz, should also be taken into account in a firm system design.) It can be seen that a data rate of 10 Mbps can be supported to a range of nearly 500 miles. For comparison, performance with a dish of one half the size (4.5 inches) at either end is given which follows the relationship that the data rate will vary as D^4 , hence the data rate is reduced by a factor of 16 relative to the 9-inch dish. More generally the data rate will vary approximately as $\text{Power}(f_{\text{GHz}})^2 D^4 / \text{range}^2$, which shows that for a given aperture size higher frequencies provide greater capacity because the transmit beam is more focused. Equivalently, the capacity will be proportional to $\text{Power}(D^2 / \text{beamwidth}^2 / \text{range}^2)$ which shows that

narrow beamwidths are key to establishing high rates; this is most easily done at microwave frequencies above several GHz.

AIR NODE-AIR NODE				
TRANSMITTER POWER	14.0 DBW	25 WATTS	10 GHZ	
COUPLING LOSS	-3.0 DB	-3 DB	0.03 $m = \lambda$	
TX ANTENNA GAIN	25.0 DB	0.75 FEET	9.33 DEG	
			49 MILE SPOT	
ERP	36.0 DBW			
PATH LOSS	-167.0 DB	300 MILES		
RX ANTENNA GAIN	25.0 DB	0.75 FEET	9.33 DEG	
			49 MILE SPOT	
RCV'D POWER	-106.1 DBW			
NOISE DENSITY	-198.6 DBW/HZ	1000 DEGREES, K	-5.0 DB G/T	
Pr/N0	92.5 DB	0 SNR BANDWIDTH=	1795.0 MHz	
LINK ATTENUATION	-4.8 DB	4.8 DB	0.01 DB/KM	
DESIGN MARGIN	-5.0 DB	5 DB		
(EB/N0)req	-6.0 DB	6 DB		
DATA RATE	76.7 DB-BPS			
	= 4.72E+07 BPS			
	= 47.21 MBPS	10. GHz		

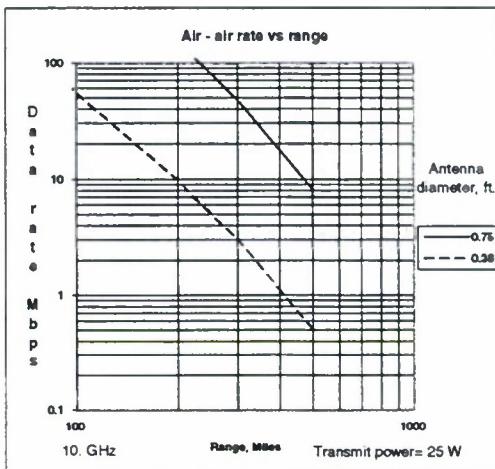


Figure 3-14. Representative narrow beam cross-link performance (Mbps vs. range).

An additional advantage of narrow beams is their spatial discrimination against out-of-beam jamming or co-channel interference. Frequency re-use for data communication should be possible. The use of spread-spectrum modulation on these links would aid both in jamming protection and more effective utilization of the spectrum through re-use.

An important design consideration will be that of spatial acquisition and tracking. Initial acquisition can be aided by the use of auxiliary omnidirectional beams as discussed in the next section. Tracking to within about a tenth of a beamwidth will be needed and can be achieved through a combination of standard angle tracking techniques (monopulse or conical scan), possibly aided by the platform's inertial navigation system.

A continuous measure of link quality, e.g., signal-to-noise ratio or missed data segments, will allow adaptation of the data rate. In turn, this can be used by the data routing algorithms to make both routing and topological decisions.

Each node will need a minimum of two beams for connectivity. As noted in Figure 3-13, multiple beams can be implemented either with multiple discrete antennas, e.g., small dishes, or with a multi-beam phased array. Although the multi-beam array approach is more flexible and capable, these clearly have to be traded-off against its greater complexity. The selection of the best technology has not been done as part of this study.

3.4.3.2 Omnidirectional Beams

The previous section showed that cross-link data rates of interest (several Mbps to 10s of Mbps) can be implemented with relatively small antennas and low RF power. However, the trade-off in taking advantage of this capability is that fairly narrow beamwidths (less than about 10 degrees) are required. (A beam with a width of 10 degrees would form a cone of connectivity of about 35 miles at a range of 200 miles.) Although it may be possible in principle for all airborne nodes to be made aware of the position of all others so that narrow beams can be pointed, it would seem prudent to provide a more fail-safe method of node discovery. Hence consideration should be given to including omnidirectional directional antenna coverage as well. This capability would also be useful in maintaining connectivity during maneuvers by the airborne platform which could result in temporary blockage of the main cross-link signal.

The omnidirectional beams could be used, for example, to announce and discover the presence of a new node in the constellation and to learn its position (via telemetry). After discovery, data can be exchanged to help make routing and topological decisions, e.g., which narrow beam on which platform should be used to point to the new node. The discovery beams themselves can also be used as a low rate backbone for data and control information. Because the omnidirectional beams will cover wide areas, they must use some form of frequency management or spread spectrum to prevent mutual interference.

The axis with the greatest degree of pointing uncertainty will be azimuth. Hence “omnidirectional” as used here refers more specifically to providing coverage of the horizon. This can be done in a number of ways including a relative narrow scanning beam which constantly rotates or a “pancake”-shaped antenna pattern which can be narrow in elevation but broad in azimuth. (It cannot be so narrow in elevation that aircraft banking would cause the pattern to miss the horizon.) The pancake pattern can be implemented with a simple vertical radiating element. Figure 3-15 shows the result of a link calculation using a vertical antenna assumed to have a uniform excitation, hence forming a pancake beam with a vertical beamwidth of $70/(D f_{GHz})$ degrees and omnidirectional in azimuth. Other link parameters are as in Figure 3-14. Interestingly, the capacity will be independent of operating frequency (except for link attenuation effects) because the gain of the transmit antenna and capture area of the receive antennas vary in exactly opposite ways with frequency. The figure shows that even with a fairly small antenna (less than 2 feet) connectivity can be established for ranges to 500 miles with data rates of about 100 kbps.

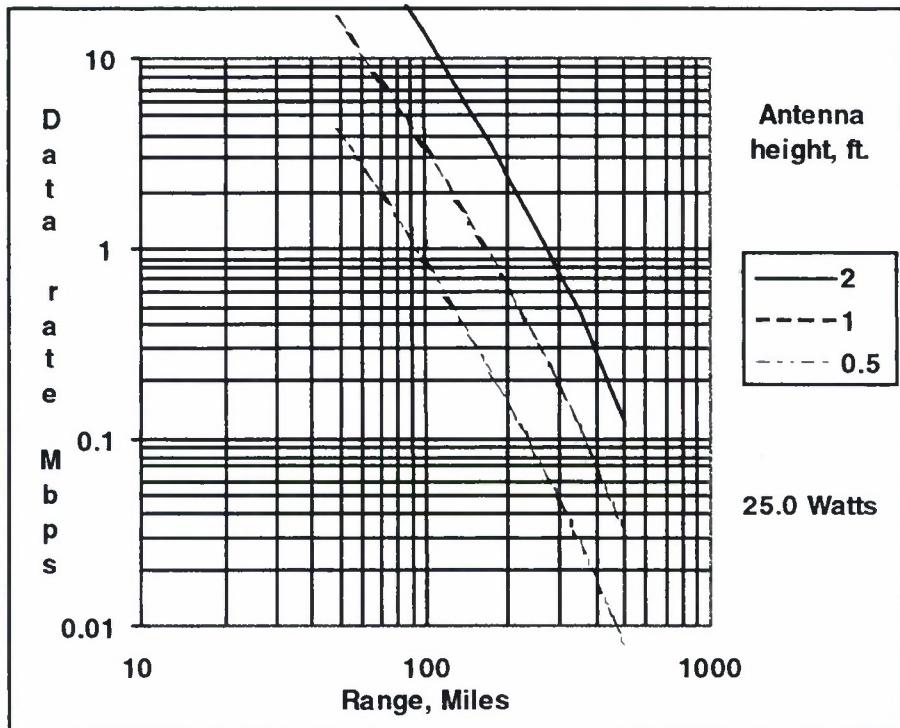


Figure 3-15. Performance of pancake cross-link antenna patterns (Mbps vs. range).

3.4.3.3 Topology Considerations

Data routing in a static network configuration can be quite complex. For the Warfighter's Internet, routing is made even more complicated by the dynamic nature of the scenarios because: 1) airborne nodes enter and leave the constellation and 2) link conditions between pairs of nodes change due to range changes, link attenuation, or maneuvers. In order to deal with this dynamism, not only will routing algorithms have to be dynamic, but the topology of connectivity should be responsive as well. There is extensive literature of general principles and algorithms which can be used as a starting point*, but this is still an active area of research. It is likely that a real-time iterative approach will be needed in which trial topologies and flows will need to be evaluated during system operation.

To illustrate the topological issues, consider the case of only four airborne nodes. If each node has only two cross-link antenna beams, then the four nodes can be connected in three ways as shown in Figure 3-16. In general N nodes can be connected in $(N-1)!/2$ ways if each node has only two beams.

* For example Bertsekas and Gallager, Data Networks 2nd Edition, 1992, Chapter 5.

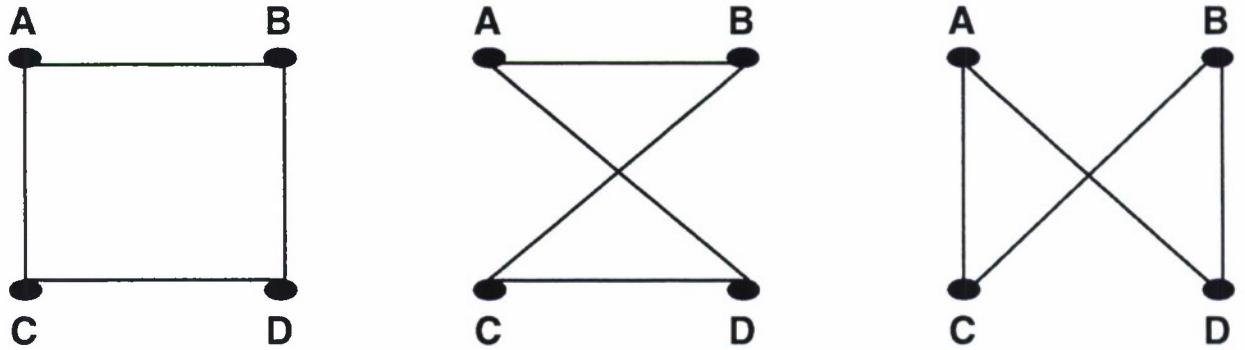


Figure 3-16. The three ways to connect four nodes with two beams each.

In a static configuration, two antenna beams per node are clearly enough to provide full interconnectivity. Two are also enough to provide interconnectivity even if one link fails as long as the original topology consists of a closed circuit. And two beams are enough to allow an additional node to join the network without disconnecting any portion.

However, two are not enough to allow reconfiguration. Reconfiguration requires at least two links to be broken if there are only two beams per node, and this would break the network into two non-connected parts.

Allowing for reconfiguration does not, however, imply that all nodes need to have more than two beams. For example, suppose only Node A in a four-node constellation has three antenna beams. Then, for example, we can reconfigure from the first configuration above to the second in two steps as shown in Figure 3-17. During the transitions to and through each step it can be verified that full connectivity is maintained (although not necessarily full flow capacity). It turns out that all reconfigurations for the four nodes can be accomplished similarly. For an arbitrary number of nodes, the smallest number of nodes requiring more than two beams in order to accomplish arbitrary reconfiguration without segmenting the network is not known at this time.

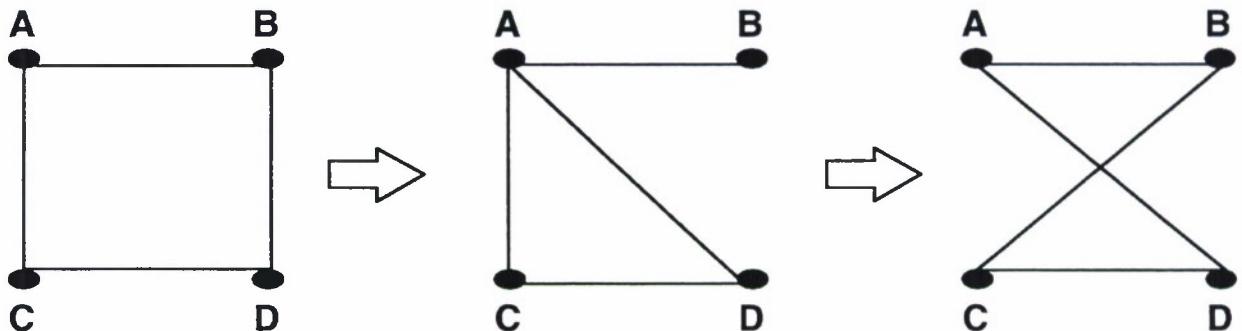


Figure 3-17. Two-step reconfiguration.

3.5 FUNCTIONAL ARCHITECTURE FOR THE MOBILE WI

3.5.1 General Architecture and Design Goals

There are different ways to implement the desired functionality of the Warfighter's Internet; each have their advantages and disadvantages; alternatives designs will be presented in Sec. 3.5.2. In evaluating the alternative designs, the differences will show up in different ways. But before the alternative designs are examined, it is first appropriate to list the design goals that any architecture should attempt to meet. And before describing the design goals, it is appropriate to set down (once again) the top-level description of the WI architecture.

At a top level, the WI is composed of a backbone (links among airborne nodes and entry nodes that are gateways to other networks) and of subscriber links (links between airborne nodes and subscribers), some of which are individual end users (with Mobile Communication Devices, MCDs), and some of which are concentrator nodes (such as Navy ships, command centers, or possibly Army RAPs).

The following paragraphs provide more detail on what is part of the mobile WI and where the interfaces are to other subnets. The mobile WI is defined as a particular subnet in a much larger tactical network. This is important, since many of the unique features of the mobile subnet can be limited to this subnet, and this subnet can be categorized as an Autonomous System (AS) which has specific properties when applied to IP networking and can fully interconnect with the external IP networks.

The subscribers, the backbone network, the entry node with support services, and connections to the non-WI infrastructure (represented here by an interface to other terrestrial networks such as MSE and DISN) are shown in Figure 3-18. The terrestrial nets are not part of this WI architecture. The interface is, however, the main way for other networks to access the users attached to the mobile WI.

A key point in this figure is that not all WI nodes are airborne. There is a required entry node which is shown with both mobility support databases and tactical database servers. This support node is connected via a standard backbone connection, either directly from an airborne node or by relay over a satellite. This entry node is mandatory, so that one would also configure alternate entry nodes to take over in the event of failure or if the nearest airborne node goes beyond direct communications range. Note that the entry node need not be located in the tactical theater; it could be on a ship, or it could be reached through a satellite relay from one of the airborne nodes. By this construct, it is not necessary to move such an entry node into a theater to initiate use of the mobile WI; this is important for support in early entry scenarios where there is no in-theater command center as yet. The entry node is discussed further in Section 3.5.3.

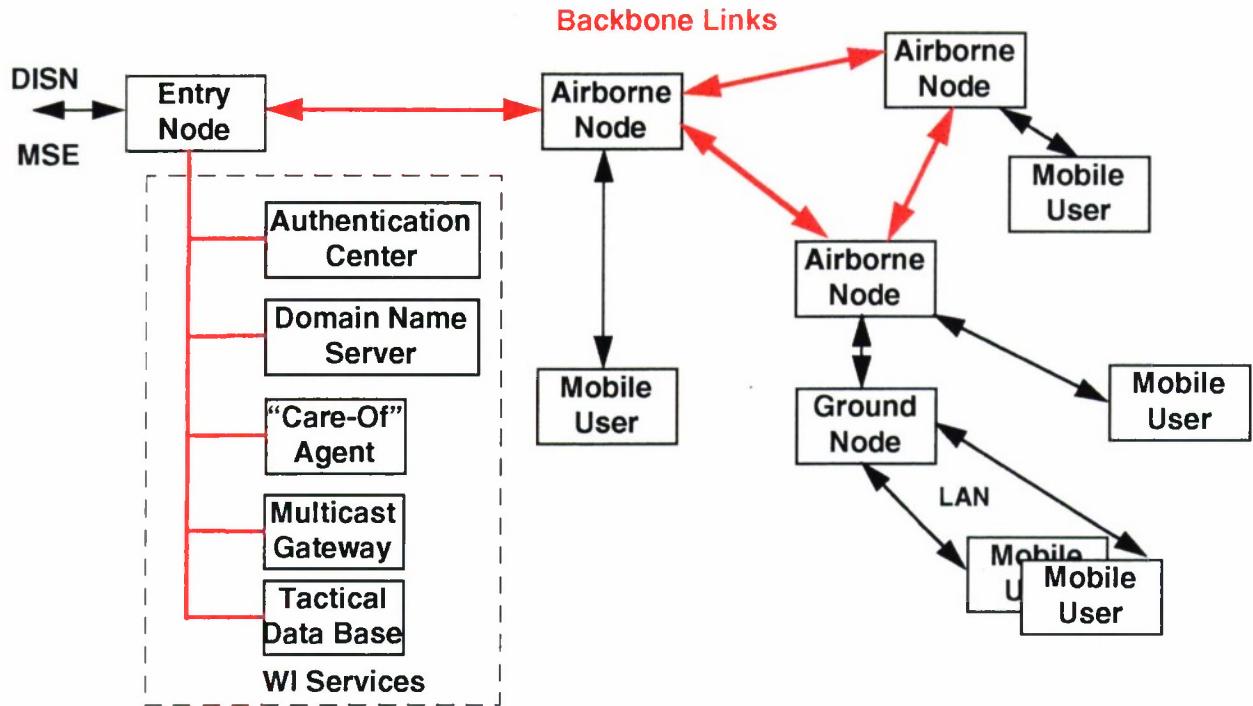


Figure 3-18. General architecture of the mobile Warfighter's Internet.

Following this general description of the Warfighter's Internet architecture, it is time to describe the design goals for the WI:

- 1) The focus has to be on providing a wide range of information services to the untethered user who can be either fixed or on-the-move anywhere where there is coverage provided by an airborne platform. The utilization of the network (information architecture) has a strong impact on the user-perceived throughput of the WI; this requires assessment of the applications and their interaction/demands on the network, and the use of smart agents, etc.
- 2) The WI must pass traffic of various kinds; specifically, both bursty datagram traffic (generated by computers) and streaming (continuous) traffic such as voice, video, and multimedia. Associated with the data source, the WI must allow specification of Quality of Service for the source traffic (in terms of delay, delay variation, bit error rate, etc.). However, it must be understood that QoS in a wireless environment may not have the same contractual meaning as in a wired network, since the wireless links are unpredictably variable in quality. In many fixed networks, the links operate with a relatively constant bit error rate; and quality of service is then determined by the communication resources that can be reserved at the nodes. In mobile networks that include fading channels, there is no easy way to guarantee QoS. As a consequence, one may in certain circumstances make dynamic tradeoffs that still allow the essential information to be exchanged. For example, in voice communications it is important to limit delay variation, and to accomplish this one may have to introduce additional fixed delay and on occasion lose data segments.

- 3) The WI signalling and access methods chosen must be efficient in the use of communication resources (bandwidth and power) in the face of the anticipated traffic types. The basic architecture is composed of a cellular-like topology for subscriber access with airborne base stations interconnected by dynamically-configured point-to-point and omnidirectional cross-links. The cellular organization shares some similarities with evolving commercial cellular telephony systems, but allows for connectionless traffic to efficiently share uplink bandwidth among bursty data sources.
- 4) Traffic may originate from users within the WI or from users outside the WI. The interfaces to the WI may see this traffic in several forms, among them are IP packets and ATM cells. The ATM cells may be carrying IP packets, or they may be carrying native ATM traffic. The WI must be able to deal with all of these. (This does not mean that the WI could not be all packet-based nor all cell-based; it could be mixed. The design issue is efficiency of transport and routing in the wireless signalling environment.)
- 5) The WI must provide a dynamic routing capability that is responsive to the changing connectivity of users and airborne nodes. Information on changes in connectivity must be propagated rapidly to insure delivery of traffic. Mobility includes not only changes of connectivity within the WI but also arrival of new users into the WI. The latter case requires interaction with the user's home network to insure that messages reach him in the WI, and the forwarding of such traffic from the entry node to the visiting user.
- 6) The WI must support efficient multicasting among users both inside and outside of the WI. WI users can be belong to multiple multicast groups and can filter on multicast group addresses as well as their individual addresses. New protocols will be required to implement multicasting in the dynamic routing environment of the WI.
- 7) The WI must provide features that differentiate it from commercial services; namely, special military-specific security features (including information warfare mitigation), user authentication and access control, anti-jam performance, user and data specific prioritization / preemption, and LPI where needed.
- 8) The basic IP technology that should be evaluated for use in the WI is the next generation IP (IPv6) which is intended to provide a better environment for security, multicasting, mobility, and QoS. It should be noted that security, multicasting, and mobility extensions are being proposed in the community for IPv4 as well, but are being defined as integral to IPv6. The major difference is that IPv6 is being designed to enable providing different QoS to different information streams. However, IPv6 does not provide for the mobility of the WI and its unreliable wireless links.
- 9) The overall communications, information, and security architectures should have a minimum infrastructure; they should be self-configuring and self-managed insofar as possible.
- 10) The WI system is not circuit switched, but relies on more efficient statistical multiplexing technologies. Circuit switched technologies normally reserve equal capacity in both directions. Half duplex circuit-switched operation is only a partial response to matching real bi-directional

information transfer needs efficiently (as in voice traffic). Pure packet switching and other options combining packet switching with cell-based technologies (as in ATM or ATM-like) are potentially much more efficient. These options allow for a considerable asymmetry in information exchange between communicating entities. For example, in a typical client-server operation, the client to server query can be many orders of magnitude smaller than the server-to-client response. In practice, a Warfighter Mobile Communication Device (MCD) can use a small portion of the uplink and still receive a fast, high capacity downlink response.

This general architecture supports the client-server information model. A user can either be a subscriber client or a server device. A Warfighter's MCD can be used for client access or an MCD can be connected to a host workstation and thus act as a server. Such a server would contend with other users on the subscriber uplinks. A tactical server can also be collocated with the entry node supporting the mobility databases. This server would be reachable via a backbone link and would not use the subscriber access uplink channels. Finally, servers can be reached via special links such as the RAP, as will be covered later.

The Warfighter MCD supports conventional voice service at a low rate, but with acceptable quality. It is anticipated that the voice rates would fall between 2.4 and 4.8 kbps; anything higher would trade marginally higher quality versus less capacity dedicated to data services.

There may be a future requirement to be able to handle secondary video distribution. The supportable downlink rates are perhaps two orders of magnitude larger than the uplinks, so video is certainly a possibility. Since the downlink is broadcast, this means that video can simultaneously be sent to all the users with a compatible video decoder. At the present time, there is commercial work going on in supporting low bit rate video over a 28.8 kbps modem for the PSTN. One could apply a variant of this technology to the mobile WI. This would require the development of low cost video decoders that can be attached to the notebook computers; the display would be on the user notebook screen. Video distribution is a good example of the asymmetrical possibilities of the WI (no scarce uplink capacity is needed).

Low rate videoconferencing is a future possibility. The requirements are larger system level capacity and particularly increased subscriber uplink channel capacity. This would require the subscriber antenna systems to have higher gains: specifically via multiple (sectorized) beams on the airborne platform and perhaps small adaptive arrays on some specially configured user MCDs. Accommodating videoconferencing will certainly raise the system cost and one must trade off the benefits gained.

3.5.2 Warfighter's Internet Architecture (Commonalities and Variations)

The intent of this section is to propose at a top level a strawman functional communications architecture and to refrain from proceeding too far into implementation-specific designs. Hence, at this time it is imperative to leave some lower level architectural decisions open since there are a number of viable options available that need more careful evaluation. Many of the concepts being proposed have never been evaluated via analysis or simulation and it is

premature to recommend a specific approach. It will be seen that at above some level there is commonality satisfying the design goals presented in the previous section. The first part of this section will attempt to present a common architecture view that applies to all variations. This will be followed by descriptions of the architectural variations. The details of these variations are found in the appendices.

It is useful to review some of the assumptions that are integral parts of the mobile WI strawman architecture.

1. There is only a single end user addressing scheme (IP addresses) and these addresses are “permanent addresses” and independent of which specific network the user is attached to. All information types are aided in their routing by the source and destination IP addresses. (However, this does not necessarily imply that within the mobile WI all route decisions are made solely on the basis of IP addressing. Neither does this preclude some “shorthand” addressing modes internal to the WI.)
2. The basic IP mobility mechanisms will be implemented to handle mobile users as seen from outside the WI.
3. The mobile wireless WI is only one network (or subnetwork) in a larger world of tactical and strategic voice / data networks. For each of the architectural options to be discussed, it is mandatory to detail the interworking issues and provide potential solutions for these issues. Similarly, the fixed global Internet is changing and it is necessary to be able to adapt to technical advances exhibited on this highly visible network.

3.5.2.1 Overview

Functional Description of the Airborne Node

Figure 3-19 shows a functional partitioning of the airborne node. While a number of functional boxes could well be combined into a single physical package, for descriptive purposes it is better to retain the functional description so that analogies can be drawn to “similar” functions being implemented in other communications systems. But similar in functional concept only, since few of the WI systems have a “usable” equivalent COTS realization at a subsystem level.

The management of the subscriber subsystem is separate from that of the backbone subsystem. The subscriber subsystem is first described and then the backbone subsystem. The subscriber subsystem contains a rough equivalent of a subscriber communications management function in a mobile telephony base station with a few important exceptions which are covered separately.

One main design goal of the subscriber subsystem is to hide from the backbone unnecessary detail on the connectivity changes of the mobile subscriber due to reassignment between (possible) multiple beams on a single platform or transitions between signalling and

assigned data paths, or in some options, between demodulators. This simplifying strategy is also routinely performed on analogous mobile telephony systems.

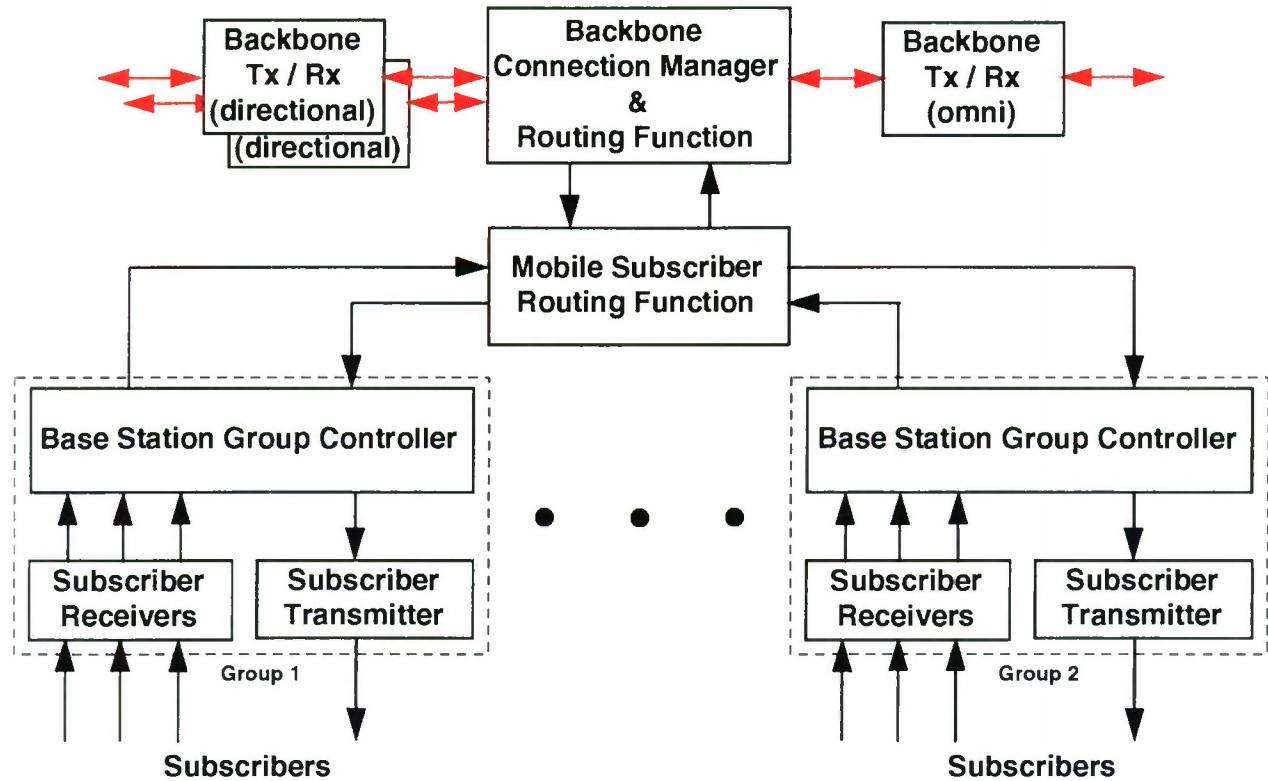


Figure 3-19. Functional view of the mobile WI airborne node communications system.

Since the airborne platform may ultimately use a multibeam subscriber antenna array to provide multiple cells, this is indicated by dividing the receivers and transmitters into **groups**. (Another interpretation of **groups** is that they serve different classes of users with different data rates or access techniques: individual users vs. concentrator nodes, for example.) This allocation is shown as fixed (as a possible realization), but it can be very flexible. There is at least one signalling channel in each beam; a user communicates over this (virtual) channel to register and be authenticated before receiving a traffic resource assignment. (There will be a number of other shared virtual channels: those for frequency and time synchronization recovery and a channel for paging/call alerting). In this example, tens of **subscriber receivers** are assumed, but because there can be many more active users than that, a “fair” and efficient demodulator sharing by the user community will be an integral part of the architecture. For each group, only a single downlink **subscriber transmitter** is shown, although this is certainly not the only possibility. The downlink stream can efficiently support broadcast and multicast in addition to the standard unicast transmissions and in-band signalling. In the figure above, the **base station group controller** forms the downlink data stream using QoS information (including priority) to distribute downlink resources. It also monitors the uplink receivers for service request messages, and puts responses of service assignments into the downlink. When subscribers first register and are authenticated,

their active presence is noted in the airborne node routing tables, which are available to the entire WI.

The active subscriber list is needed by the **mobile subscriber routing function** which performs the dynamic subscriber end link routing. It directs base station received traffic to base station downlink or to the backbone system. It also directs traffic from the backbone to the appropriate base station downlink transmitter. Another purpose of the subscriber routing function is to hide the impact of either user subscriber channel reassessments or user moving from beam to beam on the same platform from the operation of the backbone routers. For example, suppose that an airborne node is supporting multiple subscriber beams and that it is flying a race course pattern. In this case, it is normal for a ground user to continually switch coverage from one beam to another such that the end-to-end path for a communications session is continually changing. One would not want this end link route change to propagate throughout the network. Therefore, one would restrict this end link management change to the subscriber routing function and have the backbone routing functions be unaware of this local change. Only when the user changes cell coverage from one UAV to another should the backbone routing function be alerted. This is similar to the end link route management performed in cellular telephony systems.

The backbone system can also be implemented in a variety of ways. In the figure we referred to a **backbone routing function**. This term is intentionally ambiguous so as to leave room for proposing different architecture options and implementations. In all cases, from a backbone routing perspective, one only cares if the subscriber is connected to a particular airborne node, and data are routed accordingly via the backbone routing function boxes as far as the subscriber routing function which then handles the last end-link connectivity mapping.

Unlike other mobile communications systems, the backbone is also highly mobile and nodes can enter and exit at will or even unexpectedly. There is also no fixed relationship between the relative positions of the nodes. Therefore, a **backbone connection manager** is provided which uses omnidirectional cross-link communications to identify potential directed narrow beam cross-link data traffic channels. When one has N active airborne (and ground) nodes, then each node has potentially up to $N-1$ cross-links. However, a node can usually only support a lesser number of cross-links because of availability of communications assets (i.e., cross-link antennas and modems). Some of the links are non-supportable because of ground blockage and the distances may be such that the link budget could not be satisfied. One function of the backbone connection manager is to service newly arriving nodes which are to be added to the network, delete nodes that have left the network, and to assign resources to those node pairs that provide an appropriate network connectivity. The decision on the topology of the backbone system is autonomous, coordinated by the **backbone connection manager** with its counterparts on other airborne nodes. Having established what is the desired connection topology, the backbone connection manager then directs the pointing and tracking of the backbone narrow beam antennas. Once a high-speed connection is set up to a neighboring airborne node, this routing information is entered into the routing tables used by the backbone routing function; it is not necessary for the backbone routing function to “discover” the presence of the connection by use of Internet-standard protocols (OSPF, etc.). Use of the omnidirectional cross-link for traffic routing is also possible, although it is limited in capacity.

At this point, the actual routing of the traffic is handled by a routing protocol which only cares about the active connectivity. The protocol does not normally provide any information that can be used to reconfigure the active connections. Connectivity considerations and measured traffic flow are initially uncoupled, although one could provide mechanisms to change this.

Functional Description of the Subscriber Terminal (MCD)

Clearly, a critical subsystem in the mobile WI is the Warfighter's terminal or Mobile Communications Device (MCD). The remainder of the subsystems exist only to make the Warfighter's MCD a tactically effective device. The design of the subscriber MCD is particular challenging for a number of reasons.

First of all, it is anticipated that, although integrated, it will be made up of two major parts, one of which contains all the WI unique functionality and the other part of which is designed with absolutely no knowledge of the mobile WI. This does not mean that the latter cannot be integrated with the WI-unique device but rather that any adaptation should be done preferably with minor software add-ons and perhaps with already developed PCMCIA cards. The latter system would be realized as either a subnotebook computer or an advanced PDA. The advanced PDA may, in fact, be the most widely used data access device because of its extreme portability, low power requirements, and resistance to environmental conditions.

A second reason is cost. The cost will be associated with the mobile WI unique part which contains the voice access device and the radio functions. If this cost is not low (in the few thousands of dollars each range, including all security devices), then this system will not be acquired in quantity. It is only with quantity buys that one can justify the high non-recurring engineering development associated with miniaturization. With potentially thousands of subscriber MCDs, the MCD costs will dominate the entire system costs and every effort must be made to keep the per-unit costs down.

A third reason is design extensibility. If the system is to have a significant lifetime (which it will have for the military applications as opposed to commercial applications), then some flexibility needs to be built into the system. For example, assume that voice encoding scheme #1 is accepted which has a certain MOS (mean opinion score). Later, new developments produce an enhanced voice-encoding scheme that requires, say, 50% of the bandwidth but retains the same MOS. One ideally would like to download this changed scheme into every MCD and into any transcoder that interfaces to other encoding schemes. The downloading (reprogramming) could be in any combination of EPROMs and FPGAs. (The design of the airborne cellular system should accommodate any changes in isochronous data sources by avoiding fixed boundaries, such as would be associated with modest TDMA schemes.)

A fourth reason is a combination of size, weight, and power. Ideally, one would like to package the combined MCD in a fold-out water-resistant flexible enclosure about 6" x 9" x 3". The antenna may extend slightly beyond and be capable of receiving call alerts, pages, etc. Small size with high impact plastic shells should provide a reasonable degree of survivability.

A fifth reason is power availability and power management. Many Warfighters will be untethered and have to rely on lightweight batteries (C cells and lighter). The aggregate battery weight for a mission of about one week should be less than the MCD weight. For this reason, the MCD design should incorporate good active power management. Fortunately, this technology is maturing in the commercial mobile telephony and the notebook / PDA designs. In both cases, the discharge time of their associated batteries is a key competitive issue and often an important operational consideration. For cellular phones, the time between recharges assuming about a 1 hour talk time should be no less than 24 hours. In the case of portable computers, the desired battery-only operation time should be on the order of 4 to 5 hours (2 hours is probably what is often realized). Normal operations with PDAs should extend to at least a week before battery replacement or recharge.

A sixth reason is security. Unlike a commercial handset, the penalty for illegal use is not lost revenues but a compromise of tactical operations and Warfighter safety. The security aspects alone make this user MCD very unique.

There is no questioning that an MCD design with the above attributes is difficult. There are many individuals and government organizations that believe that the commercial world is best suited to perform this design. This would be true if the commercial world clearly saw such a development as resulting in an almost immediate revenue stream. The commercial world is increasingly adverse to developing technologies that could make money only 5 or more years out. There is also a false sense of confidence on how advanced our handset technologies are and particularly how adaptable the current circuit-based cellular handset designs are with respect to the necessities of the mobile WI. To make this mismatch clearer, the key WI unique requirements will be detailed both below and in the appendices.

Figure 3-20 shows the various elements making up a complete Warfighter MCD.

It is assumed that the MCD could be (optionally) integrated with a small GPS receiver. The GPS functionality can today be packaged in a volume close to a cigarette pack or less. With the presence of a GPS function, this MCD can provide useful supplementary services but we need to carefully assess whether to make these services mandatory. The main drawback to adding mandatory GPS functionality is the vulnerability of GPS to jamming.

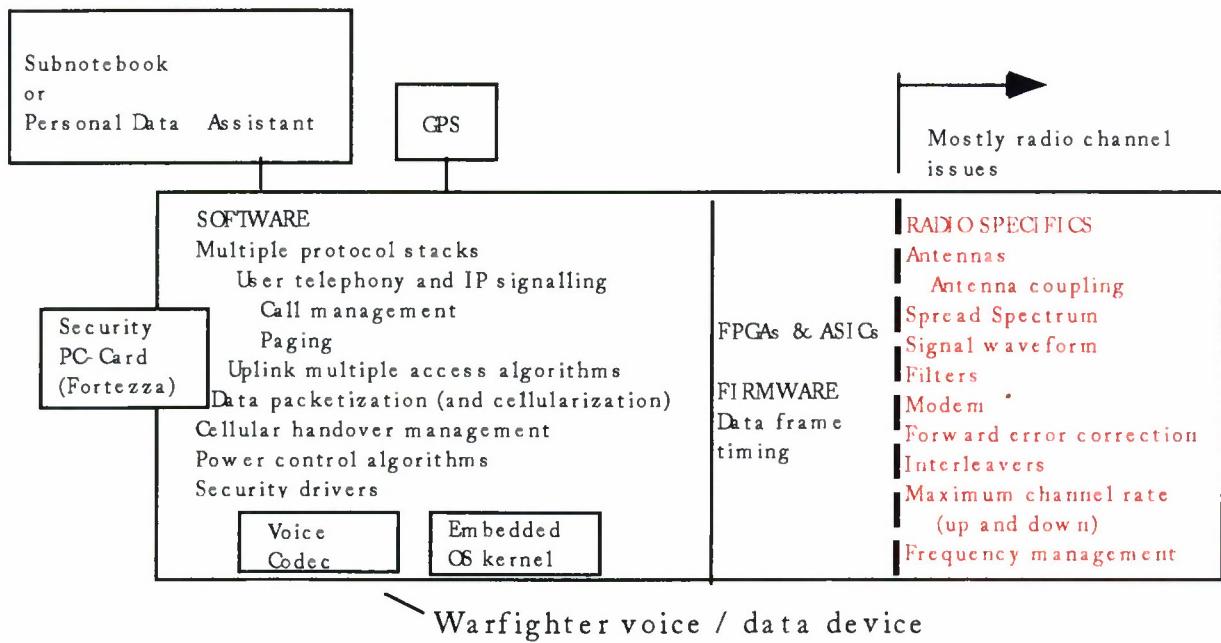


Figure 3-20. Overview of Warfighter MCD.

The functionality of the subnotebook or PDA component is not outlined although these are clearly the main user access devices for data applications and data I/O. One would expect these devices to be commercially available with any unique features embedded in PC-Cards (the new name for PCMCIA cards). The connection between the data portion and the voice/radio portion could be via a standard serial port. The subnotebook or PDA need not be active to receive (and perhaps place) voice calls and to receive pages. A small scrollable LCD display may be provided on the voice/radio device to display a page. Alternatively, an alert indicator may signal the arrival of a page message which prompts the user to activate the subnotebook or PDA to read the message. It should be observed that the long interval between recharges allows the PDA to remain on and thus be potentially more responsive in a number of operations.

The most compact MCD would include a very capable PDA. This is very suitable for adverse environments. Conventional subnotebook keyboards, mice or other pointing devices, and hard disks may not survive the elements (rain, snow, dust, sand). A modern PDA GUI could provide access to a sufficient set of tactical applications. For example, browsers are being implemented with the touch-sensitive Newton screen. Characters can be entered via an emulated keyboard or via pen input with character recognition software.

A higher end MCD would substitute a subnotebook or notebook computer. This would need more power in terms of batteries and would typically operate in more benign environments. The larger displays and keyboard would give the user more flexibility and perhaps allow applications that are not feasible with the PDA.

Security is assumed to be provided by a PC-Card Security Card that meets military security standards. The device to be evaluated first will be Fortezza, and this evaluation will include not only the usual end-to end-security encryption for user data but support for many of the other security services.

It is understandable that there is a reluctance to initiate a completely new radio design because of the perception that the military already has “too many radios.” However, with some thought it becomes apparent that the subscriber MCD has many attributes of a cellular-based phone and the military has no close analog. In addition, the airborne base station has no military analog either and it differs in many details from the subscriber MCD so it would require a new development as well. The asymmetry in complexity between the subscriber MCD and the airborne base station are intentional since additional functionality is normally added to the base station to simplify the design and lower the cost of the subscriber MCD. We have covered elsewhere the great asymmetry between uplinks and downlinks (data rates, multiple access, channel sharing, etc.) and these need not be repeated. Note that most terrestrial radios (land mobile radios and packet radios) are designed to interoperate in a symmetric fashion. It is only when there are hub or cellular-based radio networks that asymmetry is encouraged and a necessary part of the design. This does not mean that an entirely new development effort is necessary, but rather that a careful design will allow the integration of a number of existing technologies and available sub-elements (both commercial and DoD) to create the “new radio” from these building-blocks.

In this section, the intent is to largely ignore the radio channel specific issues and the low level error correction techniques (interleaving and FEC) and to concentrate on the implementation issues that make this device considerably different from all other terminals. Generally, the radio specific issues are implemented and enforced at the physical level. Some of the radio control issues are invoked at higher protocol levels.

Thus, the focus will be on the software components and some of the firmware needed to provide WI unique support. Many of the radio channel characteristics and frequency management issues are common with some cellular systems. Perhaps the most unique characteristic would be to provide secure patterns to feed the selected spread spectrum implementation. The spread spectrum choice could also require a unique implementation, but one similar to previous systems.

In this section, the intent is not to cover all the warfighter voice/radio device functions but rather to point out where they are essentially the same or differ from analogous functions being built into commercial systems.

Antennas - small omnidirectional antennas are being fabricated inexpensively for the frequency regimes from about 800 MHz to 2.6 GHz. There are also antennas available in the two ISM bands. Antennas should not be a major design issue.

Spread Spectrum Techniques - Military spread spectrum techniques are used to provide covertness, protection from intentional jamming, and protection from mutual interference of other users. The commercial SS techniques generally rely on known sequences, not secure

sequences, for their implementation. While this is only a small conceptual change, it does have implications on acquisition, synchronization, and timing.

Frequency Management and Reuse - Compared to the commercial systems, WI is more complex. Commercial systems have pre-planned, stable frequency utilization plans, while the military cannot. The desire is for larger bandwidths to provide more AJ. Similarly, the reuse strategy is more complex since the backbone nodes are moving relatively to one another and there is little regularity or predictability in beam footprint spacing.

Cellular handover management should have many strong analogies to commercial systems except that the coverage will seldom be as uniform as in a preplanned cellular system. Other differences can be found in the handover protocols which will be closely coupled with frequency reuse and must be coordinated with the frequency reuse plan. From the user terminal perspective, handover should look quite similar to the commercial case, but from the airborne node perspective, handover will be implemented in a significantly different fashion.

Protocol Stacks - This is the area of major design differences. Most cellular systems are designed around switched voice circuits. Data is an overlay on the voice circuit and the available rate for data is thus less than the voice data rate. In the WI, the need for special signalling messages still exists although the message set is considerably different, and data may often be the primary service. More significantly, either a packet- or a packet/cell-based information packaging is used. These also include unique signalling tasks not encountered in circuit-oriented designs.

Basic MCD power control algorithms should be similar between the WI and some commercial cellular systems. However, the parameters chosen for the power control algorithms will be different from the commercial world since the channel is packet/cell-based and not circuit-based. In some cases, LPI consideration may require new algorithms for power control not seen in the commercial world.

It is anticipated that greater than 90% of the completely new work will be associated with the software aspects of the MCD. Some of the networking code will be derivatives or evolutions of TCP/IP and this class of code can be obtained and modified as needed. However, the code and algorithms for cellular beam management and handover are normally proprietary and not public (this includes all the associated signalling). This means that one would have to acquire government rights to these codes or else initiate a new design.

Interfaces Between the Airborne Node Basestation and On Board Communication Assets

The importance of interworking between the mobile WI and other communications networks has been previously stressed. These interfaces were described as occurring primarily through the Entry Node, where connection between the WI and networks such as the MSE and the DISN will take place. And while this interworking is sufficient to provide connectivity to all other networks, so long as they connect somewhere to the DISN or the MSE, this kind of connectivity may be a little more remote than desirable. As was shown in Figure 3-1, it is also

possible to have interfaces to other networks directly on board the airborne nodes hosting the WI, providing the airborne node is carrying some element of that other network. If the Airborne Communications Node (carried on board the Global Hawk UAV) also hosts the WI, there will be a number of such communications systems on board with which the WI might directly interface. Not all of the elements of the ACN are networking-oriented, however, but some are.

For example, a SINCGARS relay carried on the Global Hawk could be used to deliver (or receive) IP packets through an INC (InterNet Controller); the WI could receive (or deliver) these packets to the WI by treating them as if they had been received over a “backbone” connection, and would enter (leave) the backbone routing function just as does regular backbone traffic. This does, however, raise significant issues with WI resource allocation for this SINCGARS traffic, and all of the network signalling and end user mobility management this implies. (For example, the users attached to the SINCGARS on the ground should get entered into the WI registration database for future routing decisions.)

A similar, but more complex connectivity that may be desired is with the Army Tactical Internet (TI), a ground-based “packet-radio”-like network of SINCGARS and EPLRS. Connectivity between the TI and MSE is normal. The mobility of this connectivity will be enhanced in the future if the Radio Access Point (RAP) with its High Capacity Trunk Radio (HCTR) is deployed. The RAP serves as a aggregation point for the TI links, forms ATM trunks through the HCTR, and relays to other RAPs or to an MSE extension node. The RAP is a tracked vehicle, the better to keep up with the highly mobile fighting forces; however, it may travel further than Line of Sight (LOS) links to the MSE permit. For this reason, it has been proposed that the RAP might use an airborne relay, such as the ACN, to extend its range BLOS. But once the RAP ATM trunks are present in the ACN, there is the added possibility of arranging interworking between the RAP circuits and those of the WI. The TI is an IP network, but the IP packets are carried in ATM cells in the HCTR transmissions. Thus the interface to the WI on board the ACN must be able to handle this format and convert it to whatever internal format the WI uses. As with the case of the single SINCGARS radio, this raises considerable issues with WI resource allocation and all of the cross-network signalling and end user mobility management this implies. A conceptual diagram showing possible interconnectivity between these two systems is shown in Figure 3-21; a complete discussion of this possibility, some variations, and some implications is contained in Appendix K. (The figure assumes that the Individual Subscribers in the WI are using IP from handsets, the MCDs.) The blue dotted lines show the normal WI connectivity, while the green dashed lines show connectivity between RAPs that can make use of the WI cross-links (CL). The connectivity between the RAP and the Individual Subscribers (IS) is shown in red, and requires conversion between IP-over-ATM (RAP format) and IP (the example WI format) taking place on board the ACN. (The figure does not show a red path using the cross-links, but it could very well do this, too, in either IP-over-ATM or in IP format.) This figure implies a very tight and intimate connectivity between the TI and the WI, in which (for routing to take place) each system will need complete knowledge of the routing information about the other system. This is no small problem, and further study of its implications is necessary. In the meantime, the WI architecture alternatives to be described in following sections do indeed make provision for interconnectivity with “other” formats on board the airborne nodes, but say very little about the more complex interworking implications.

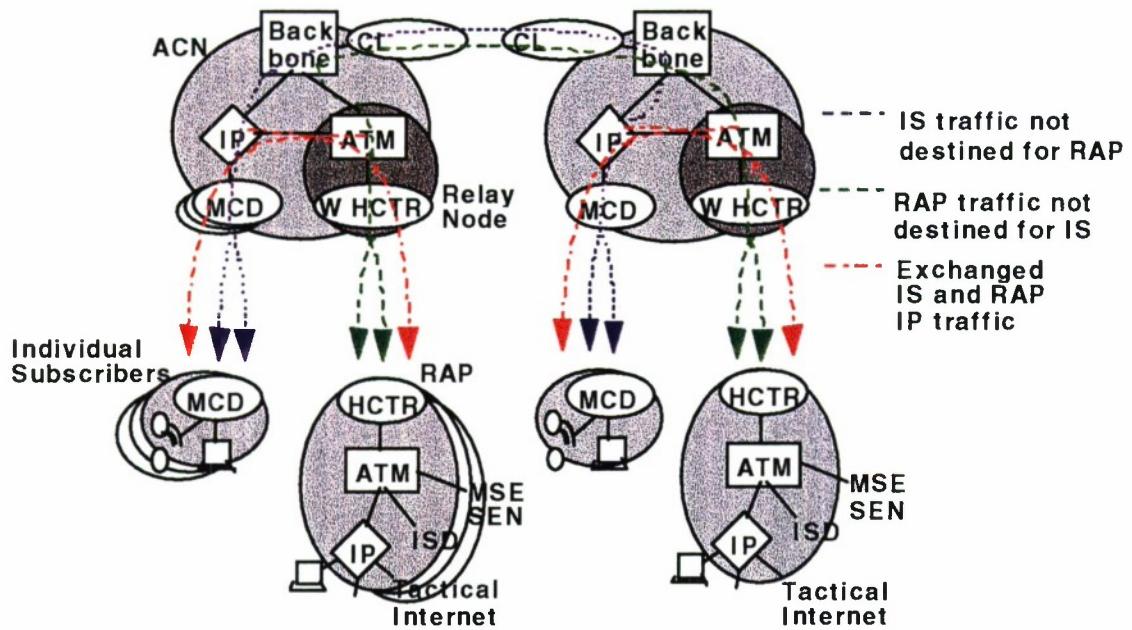


Figure 3-21. Warfighter's Internet and radio access point interconnectivity.

Introduction - Architectural Options

In the next two subsections, two different architectural options are reviewed to highlight the major differences in approaches. Subsequent tradeoff studies will undoubtedly change the details of each architecture and hopefully converge on a common approach. The details of each architecture will be placed in more comprehensive appendices. Much of this report will explore the spectrum of options and point out the relative advantages of each.

All the architectural options will address the provisioning of variable QoS. The architectural options will hinge on which entities actually control the reservation of resources that provide QoS differences. In the older IP-based systems, one could not predict QoS because of the non-restrictive compact that IP implicitly makes with its diverse underlying subnets. In an IP system based on IPv6, the routers will be designed to provide the reservation of communication resources needed to enforce a specific QoS. However, as previously noted, wireless links are unpredictably variable in quality, and no system can give absolute guarantees of a specific QoS.

Cell-based technologies (like ATM and the NRL MCA) were developed specifically to manage communication resources efficiently and to provide different levels of QoS. When IP runs on top of cell-based subnets, there is some synergy between IP and the underlying cellular subnet that cooperatively controls the resources. The cell-based options will detail the interactions between the cell-based subnets and IP. This is necessary since much of the data applications world is IP based.

3.5.2.2 Summary of Architecture Option 1

The objective of this option is to transfer all information in IP addressed packets throughout the subscriber and backbone subsystems. There will be no packet segmentation into “cells.” In terms of the generic architecture just covered, the backbone routing function is entirely assumed by a router running IPv6 protocols. The operating assumption is that IPv6 configured routers can provide the QoS levels required within the limits of the dynamics of the fading radio channels.

3.5.2.2.1 Airborne Node (Option #1)

Figure 3-22 shows a functional partitioning of the airborne node for option #1.

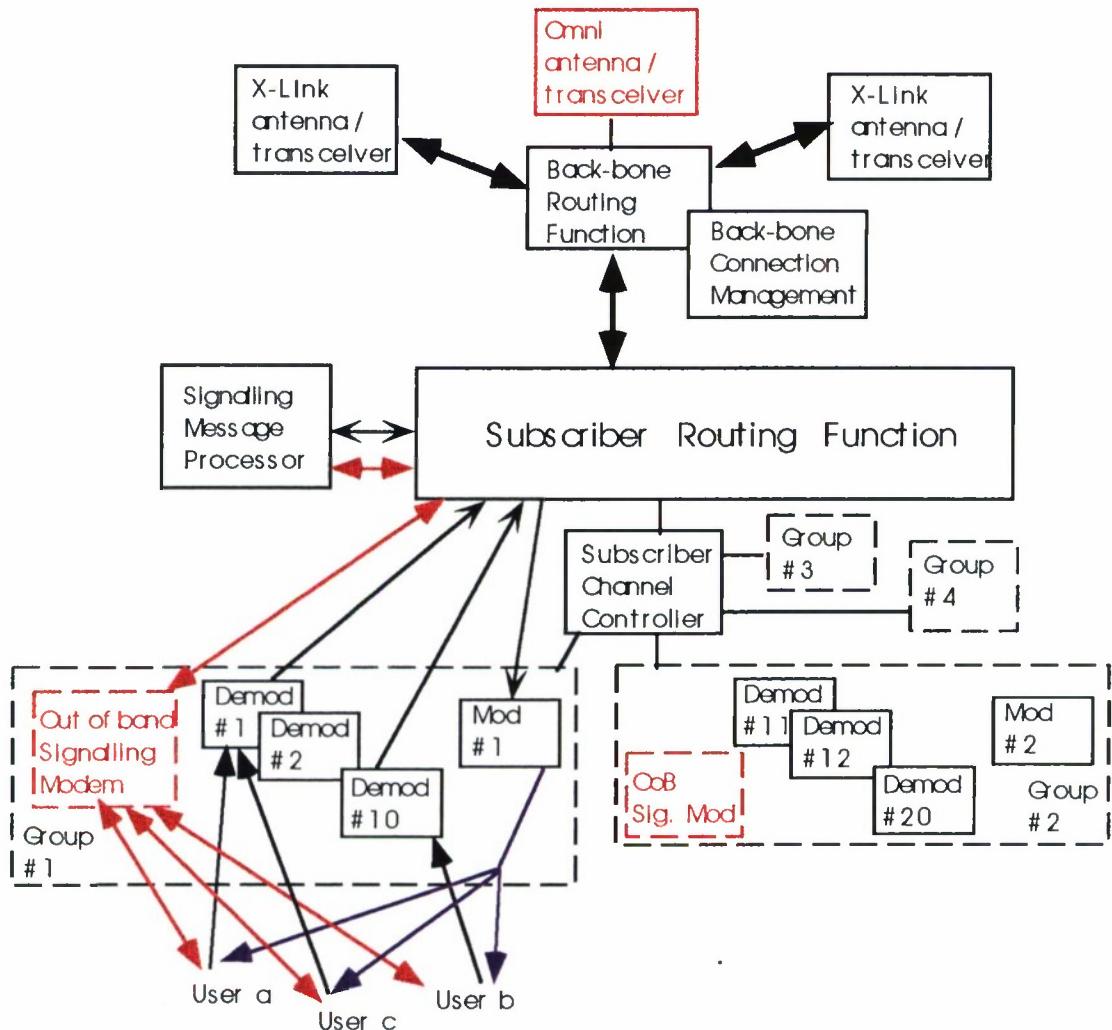


Figure 3-22. Functional view of the mobile WI airborne node communications system.

Signalling

Signalling comes in two flavors: one branch evolving from telephony and the other from data networking. Since telephony-like operations are being blended with traditional data networking operations, it is useful to use the generic term signalling to cover both cases. For data networking, the term signalling is also used to represent all the communication mechanisms and messages excluding those directly associated with carrying user traffic. This distinction will become clearer. It is also useful to show the network protocol stacks that describe pairwise interactions between peer subsystems.

The telephony signalling is used principally to establish and control access to specific communication paths so that more traditional data networking signalling can be run. This separation of signalling function is common. In many WAN applications, the user invokes telephony signalling to get to some network access point. Once this is accomplished, the data network protocols (i.e., signalling) are invoked to establish a data exchange session.

In cellular telephony, it is not uncommon to split telephony signalling into out-of-band and in-band signalling. Out-of-band signalling refers to the operations prior to the user being authenticated, not that it is necessarily done in a different frequency band. In-band signalling refers to the situation where the user may be active in an information exchange session but still needs signalling to interact with the “network.” In-band signalling implies sharing communication channels with user information. The separation is done partially to eliminate the need for the user MCD to operate two concurrent modems.

The figure shows a separate **out-of-band signalling modem** function which allows inactive users access to the network. When not active, a subscriber is normally in a passive listening mode and is periodically synchronizing to the downlink (on this out-of-band signalling modem). After downlink synchronization completion, the user listens to the downlink out-of-band signalling for appropriate messages.

To become active, an uplink synchronization process is executed and the user is then ready to use this out-of-band resource for uplink signalling. Users access this channel to register with the system before any other communications resources are authorized for use. The registration process is coupled with an authentication process that verifies the user and associates the user with specific communications resource capabilities. In this scheme, after a user has been assigned to a demodulator and modulator, these assigned assets will carry not only the user information traffic, but also any subsequent signalling associated with this user. This latter function is the in-band signalling.

The out-of-band modem function would also handle call alert and paging for inactive users. Registered users assigned to a traffic demodulator(s) and modulator will receive their call alerts and pages as in-band signalling. The alternate use of signalling on either the out-of-band or in-band signalling facilities allows the handset to use a single instance of a modem as noted earlier.

The out-of-band modem function is the first line of defense in protecting the bulk of the communications resources from security threats. By carefully controlling access to this channel, mitigation of a “denial of service” attack is possible. It is clear that potential users need to be able to listen to a signalling channel before becoming authenticated; this is provided by the out-of-band modem function. Moreover, the normal downlink broadcast remains “hidden” until the user is authenticated and has been granted the rights to listen to that broadcast. A valid security goal is to prohibit an unauthenticated user from listening to the general downlink broadcasts. Even if the user traffic is encrypted, there is the possibility of performing a traffic analysis on the broadcast downlink stream when the addresses are not hidden.

All user telephony-like signalling is handled on the airborne platform by the **signalling message processor**. For out-of-band signalling, the signalling message processor is the default connection to the out-of-band signalling modem. Once the signalling is performed in-band over the traffic channel, the messages are routed via the subscriber router between the user and the signalling message processor. The signalling message processor also communicates with other system entities that respond to commands (with parameters) implied in the signal messages.

However, at a higher protocol level there will be provision for IP, and the IPv6 support protocols are another type of signalling that is not supported by (or even passed through) the signalling message processor. Routing on an end-to-end basis (connection management) could be performed by an IP routing protocol and not handled by the signalling message processor. [This is analogous to communications over the Internet from one’s home. A dial-up session connects one by a voice-capable link to the Internet Service Provider (ISP); this is a pure telephony interaction supplemented with modems. Once this connection is established, end-to-end communications, including routing, are accomplished via standard PPP/SLIP and TCP/IP.]

To be more specific, the signalling message processor is meant to handle those messages that deal with the subscriber uplink and downlink channel assignments. The actual management of the uplink and downlink channels is performed by the **subscriber channel controller** which handles uplink and downlink channel assignments and maintains an active list of subscribers and their communication capabilities. The latter includes user priority levels and the types of information that they are allowed to exchange. For example, one may want to restrict low rate videoconferencing to only specific individuals. The reason is that even low rate video is expensive in terms of communication resources.

The priority levels and communication capabilities are relatively static entities. There is also a much more dynamic channel assignment function to be coordinated by the subscriber channel controller. We are referring specifically to the multiple access schemes implied by the sharing of the uplink demodulators. As noted earlier, a number of uplink multiple access schemes will be examined, ranging from random access schemes, to various highly dynamic reservation access schemes, to less dynamic schemes like TDMA. Within the option #1 architecture, a number of multiple-access possibilities will be evaluated.

The role of the **subscriber routing function** has already been described in the overview section on airborne node commonalities.

In a “pure” IP scheme, the **backbone routing function** would be handled entirely by a router running IPv6; there is no need for an additional switch. Again, from a backbone routing perspective, one only cares if the subscriber is connected to a particular airborne node, and data are routed accordingly via the backbone routing function boxes as far as the subscriber router which handles the last end-link connectivity mapping.

The role of the **backbone connection manager** has also been discussed in the overview section on airborne node commonalities.

3.5.2.2.2 Warfighter MCD (Option #1)

Figure 3-23 shows only a notional representation of a typical MCD. The equipment shown consists of a radio similar in size to a cellular phone with a small data input/display device. The size shown emphasizes the desire for the Warfighter to operate in a completely untethered mode and an ability to operate on the move.

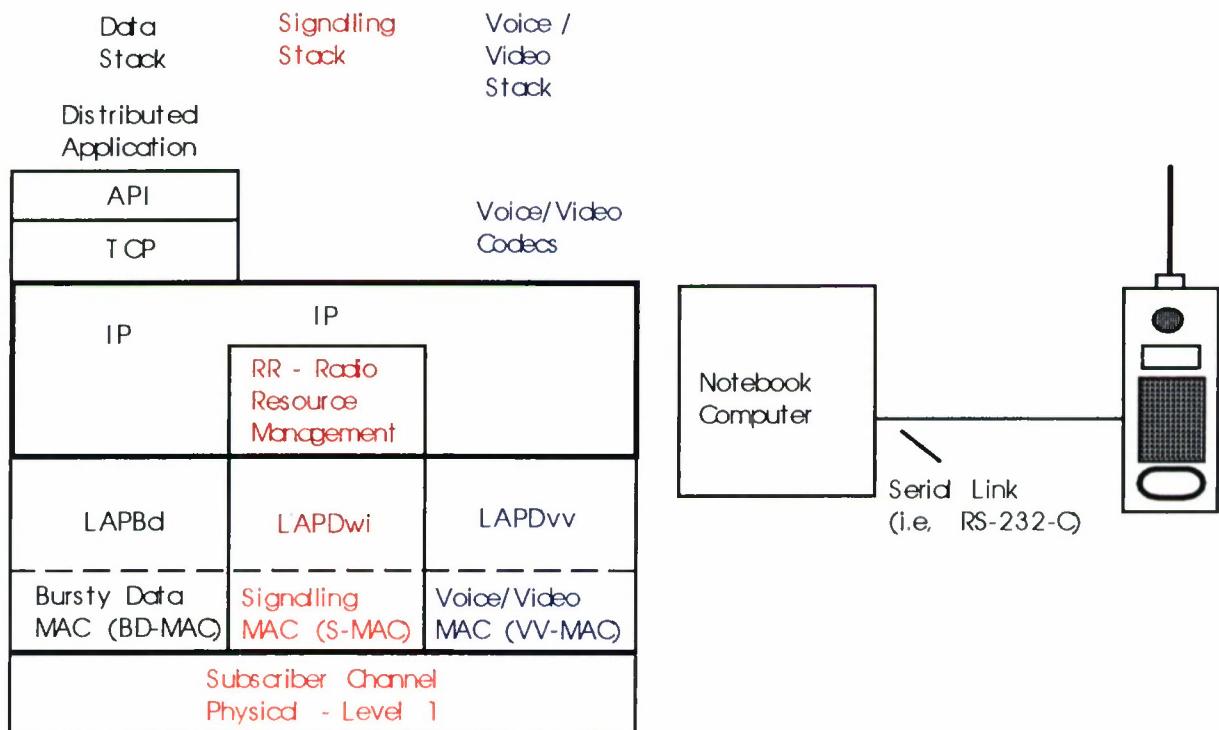


Figure 3-23. Notional architecture of a representative Warfighter MCD.

A variation on a high-end MCD would be configured to act as a gateway between a legacy terrestrial tactical IP network like SINCGARS and the mobile WI. In this case, the notebook would be executing a gateway router function. In this situation, a user would request an uplink and downlink for an indeterminate time and then transfer control to the gateway function. The

fact that there is a long-term assignment of uplink and downlink assets does not mean the channel will be used inefficiently. The channel will only be used when there is actual traffic being transferred.

MCD power management is a critical function. The average radiated power is less than or equal to about 3 Watts. Since the radio antenna is essentially omnidirectional (with a gain close to 0 dB), there is no antenna pointing required and the MCD in almost any position should be capable of receiving call alerts or pages. In the normal operational mode, the MCD is actively monitoring the paging/call alert channel on an intermittent basis, perhaps during time slots when pages/alerts are sent.

Downlink synchronization is performed periodically and depending on operational policy, full-power active uplink synchronization and registration (or re-registration) is performed periodically but less often. This registration process may provide user position to allow the planners to determine tactical asset location, provided that a GPS receiver is included. (The compact GPS subsystem is not shown in the figure. GPS is not mandatory for the operation of the system, but is a useful supplement, particularly at the application level.)

The figure also shows the types of protocols that the MCD must support. The actual partitioning of the protocols across the components is not shown here, but is detailed more in Appendix E. We include multiple protocol stacks to show that the protocol stacks are not like the single OSI stack used in tutorials. The uplink protocol stacks are used as an illustration. There is a signalling stack which controls how the network is to be managed (calls initiated, taken down, routes determined, etc.) As mentioned earlier, there are two types of signalling (telephony and IP oriented) and the signalling protocols cover both. Although we imply three completely different uplink multiple access protocols for signalling, data, and voice, this need not be the case. This was done for generality and leaves room for implementation options even within Architecture option #1.

We also have hidden protocol differences between uplink and downlink (the figure focuses on the uplink). On the uplink, the user-generated signalling messages statistically share a multiple access channel (MAC) with other users who are also initiating signalling. This protocol is named S-MAC which is analogous in concept to Ethernet MAC. Since the downlink is a coordinated broadcast there is not a MAC equivalent. Some of the signalling messages were requests for a traffic channel. The signalling messages were interpreted in the subscriber message processor, and IP datagrams were sent to the Authenticator to confirm suitability of the request. Responses from the Authenticator are relayed to the subscriber channel controller to set up the channels and also to the user to indicate the channel assignments. The Radio Resource Management protocol is the mechanism used to coordinate the assignment of uplink and downlink channels. It is noted that the Radio Resource Management function is a sublayer within the network layer.

Channel assignments are made partially on the requested information use. If data communications are the intended use, the transfers would take place over the data protocol stack at the left side of the stack. LAPBd is similar to LAPB (the index "d" implies data) and can support acknowledged and unacknowledged HDLC-like frames at the link level. The acronym

LAPB is only used to draw an analogy with the cellular signalling standards. LAPB with a suffix (d, s, or v) implies a new mobile WI protocol.

Voice and video uplinks could be handled differently from data even though both data and voice are packetized. The MAC layer could well represent a TDMA access scheme. The upper link sublayer protocol (labeled LAPDvv) would use an unacknowledged link level framing structure to carry these packets.

It was mentioned earlier that because the downlink is a shared but coordinated channel with one transmitter (the airborne node), there is no need for a conventional multiple access scheme. It is, however, a scheme which freely mixes packets destined for different users and hence there is a need for a way for the MCD to recognize (filter) packets addressed to it. One could filter on the basis of IP addresses, but this is computationally intensive; one would like to filter at a lower protocol level. One would want the equivalent of a machine address and this can be implemented so that one could have a simple address filter sublayer in the place of a MAC-type sublayer. While this is an implementation detail, it represents the type of consideration that needs to be addressed to reduce computational (and power) requirements.

The Warfighter MCD can contain one or more PCMCIA cards (now called PC Cards). This is discussed further in the section on security.

Further discussions on the Warfighter MCD will be presented in Appendix E.

3.5.2.2.3 Interfaces Between the Airborne Node Base Station and Other On Board Communication Assets (Option #1)

All options need to address connections to legacy systems and particularly to emerging tactical ATM-based systems. Just how far ATM will penetrate into the mobile WI is part of the architectural options. In option 1, the initial position is to provide a gateway between the ATM subnet and the subnet representing the mobile WI.

Figure 3-24 shows one method that provides an ATM-type interface located in the airborne node. The assumption is that an Army RAP (Radio Access Point) compatible subsystem is located in the airborne node and that a ground server is accessible via this trunk. If the desire is simply to relay between such RAP ground sites, then a repeater function on the airborne node would suffice. However, if there is the need to cross-couple such a data pipe with signals to/from individual subscribers (described above) within the airborne node, then the ATM trunked uplink would have to be broken apart and individual user sessions routed over the backbone. This would require additional on board demultiplexing, and switching/routing equipment, which could be interfaced to the WI backbone router as shown. It is noted that the broken out sessions need to be converted to IP packets. This is not a problem if one were carrying IP over ATM, but would require additional design if there were only native ATM packets. In fact this technique could be the generic interface mechanism for any other (non-WI) communications service carried on the airborne node.

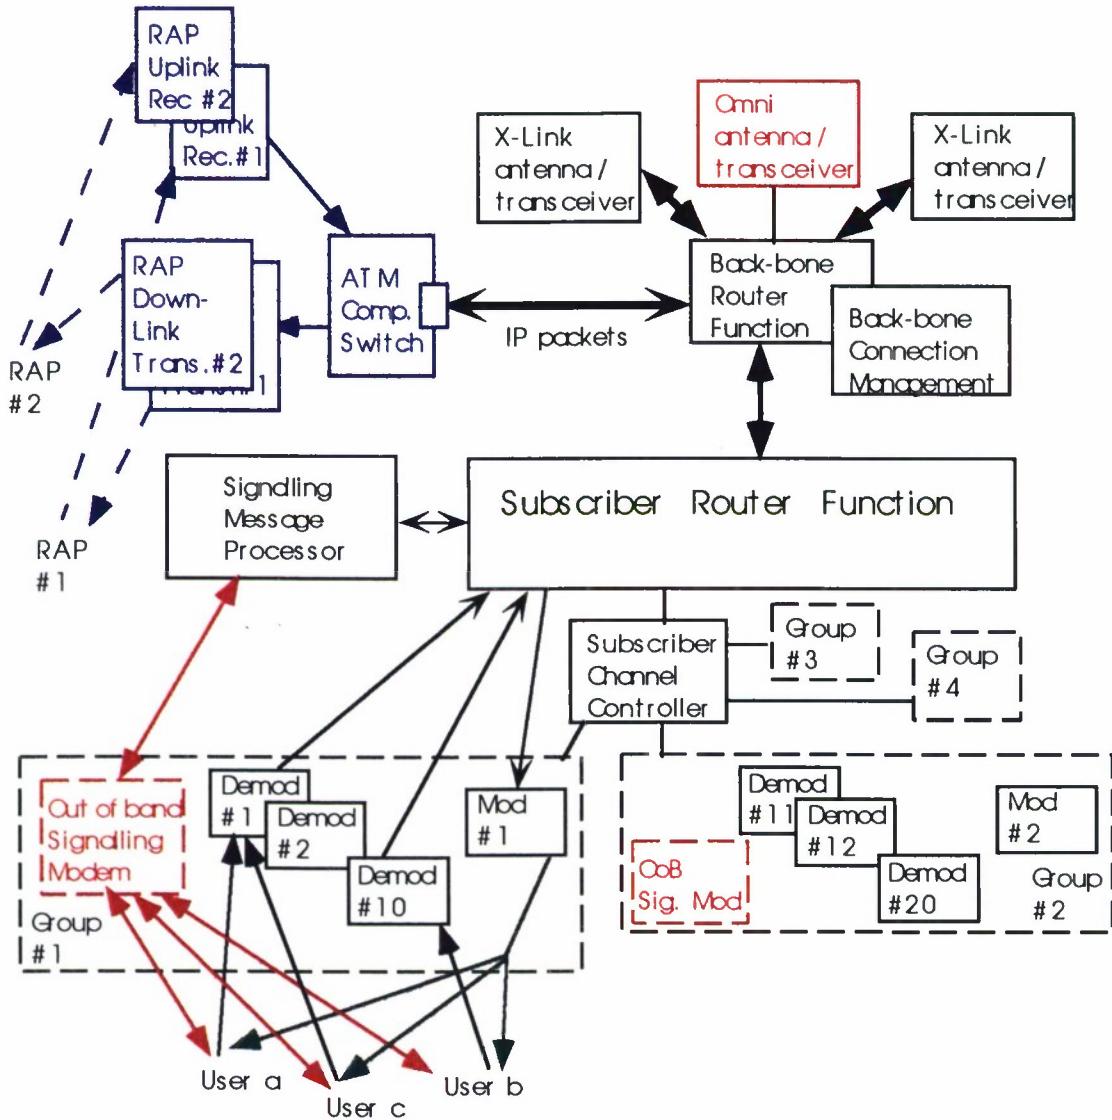


Figure 3-24. Base station interfaces to on board communications assets.

As indicated in Figure 3-24, the airborne node can simultaneously support as many ground RAPs as there are instances of the complementary airborne transceiver. With only a single RAP-compatible transceiver in the airborne node, the only function of the ground RAP would be to interface to the Backbone Router so that RAP IP packets could be transmitted either over the backbone to other airborne nodes or to directly connected subscribers under the same airborne node footprint. With only one RAP-compatible transponder, the ATM-compatible switch would only be used for converting ATM cells to IP packets; there is no ATM switching function to other RAPs needed.

If more than one RAP under the same beam needs to be serviced, then multiple instances of the airborne RAP-compatible transceivers would be needed. Placing two RAP transceivers in the airborne node would then allow direct interconnection between two RAPs on the ground. At

At this point, the ATM switch could then be used both to interconnect the two RAPs and to connect each RAP to the backbone. There are a number of issues of placing multiple RAP transceivers in the platform, such as frequency management, assignment of ground RAPs to specific on board transceivers, etc. These are being addressed in an on-going CECOM program.

3.5.2.3 Summary of Architecture Option 2

Architecture Option 2 defines the following major system components: 1) WI Backbone Cross-Links Segment, 2) WI Cellular/PCS Segment, 3) IP Router Interface Segment, and 4) ATM Switch Interface Segment. Figure 3-25 shows these components, the WI boundary, and several important interfaces.

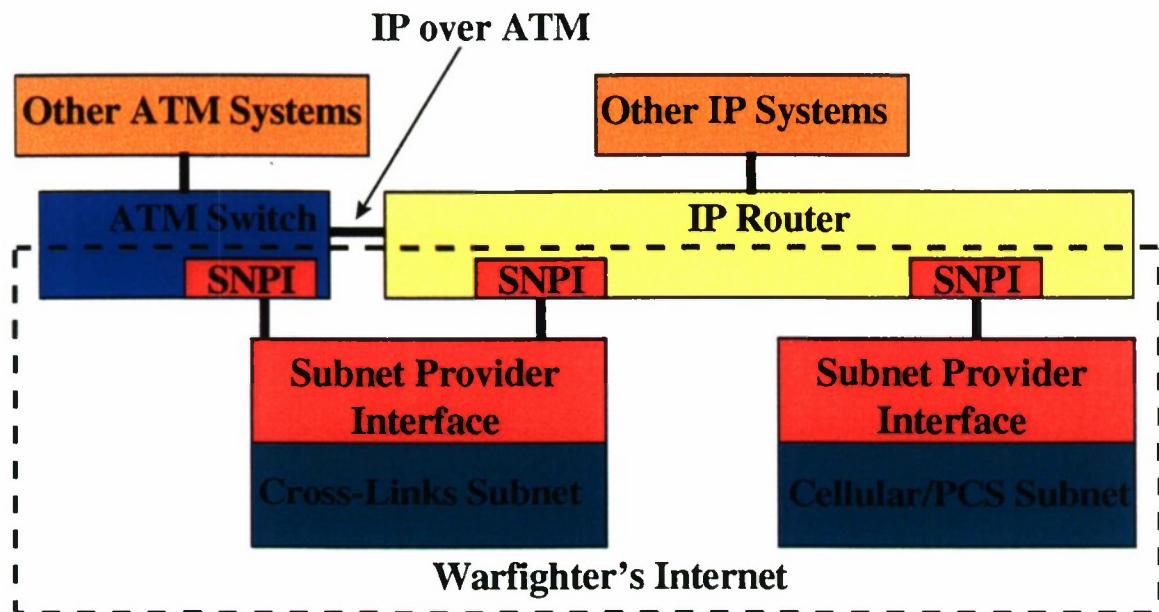


Figure 3-25. Architecture (option 2) system level context diagram showing WI components, boundary, and interfaces.

External traffic delivered to WI is assumed to be one of the following three types: IP-over-ATM, native ATM, or IP. IP-over-ATM traffic is passed to the IP router, which unbundles the IP packets and forwards them to WI. Native ATM and IP traffic tunnel through the WI subnets. External traffic is forwarded to either the Cross-Links or Cellular/PCS subnets via the Subnet Provider Interface (SNPI). The SNPI has a component (client) that resides in the ATM switch and IP router and a component (server) that resides in each subnet. The IP Router Interface (SNPI IP client) and the ATM Switch Interface (SNPI ATM client) enable WI to connect to standard IP routers and ATM switches and subsequently to external systems.

The most significant feature of Architecture Option 2 is that both the Cross-Links Segment and the Cellular/PCS Segment are designed as subnets. That is, each of these subnets is responsible for routing traffic between the nodes that make up the subnet. Figure 3-26 is an

example illustrating the nodes and links that make up each of these subnets. In subsequent sections we describe these subnets as well as the IP Router Interface and ATM Switch Interface segments of the architecture.

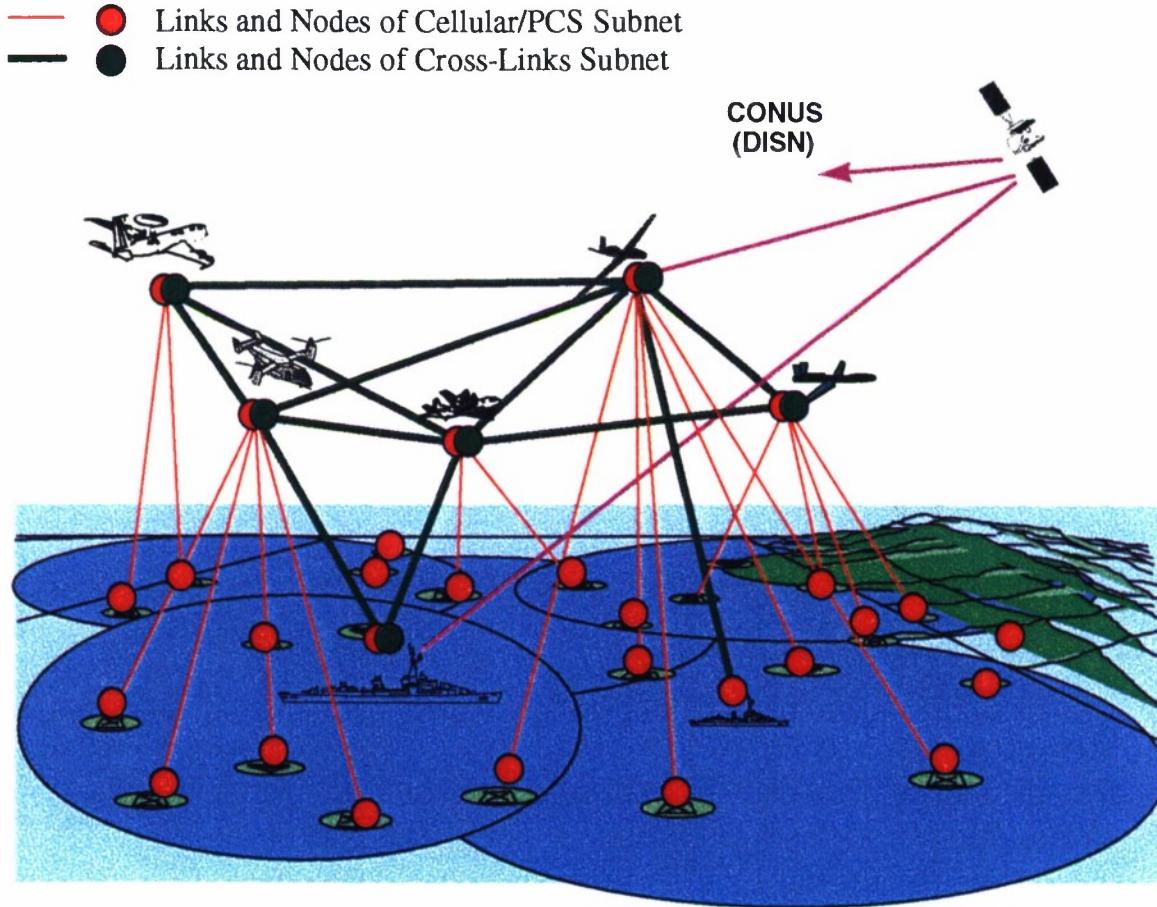


Figure 3-26. Architecture (option 2) is characterized by a Cross-Links Subnet and a Cellular/PCS Subnet, whose nodes and links are shown here.

3.5.2.3.1 WI Cross-Links Subnet

The Cross-Links Subnet serves as a Wide Area Network (WAN) for the Warfighter's Internet. In order to provide a wide area of coverage, several of its nodes are airborne. The maximum number of nodes in this subnet depends on the link data rates and several other factors that limit it to about 100.

Option 2 assumes the same suite of hardware for the WI backbone nodes as Architecture Option 1. This includes a Medium Data Rate (MDR) (100s of kbps) RF system using omnidirectional antennas and a High Data Rate (HDR) (few Mbps) RF system using directional

antennas. The omnidirectional system uses one transmitter and a bank of receivers. For Cross-Links subnets with many nodes, this bank of receivers will probably be limited to 10 or less. The number of directional antennas that will be available is unknown at this time, but it will probably be about 2 or 3 per platform.

The two architectures use the same hardware architecture; however, they differ in how this hardware is deployed. Figure 3-27 shows the architecture of an airborne node for Option 2. Note that communication over the HDR links is managed by the Cross-Links (X-Links) Subnet in Option 2.

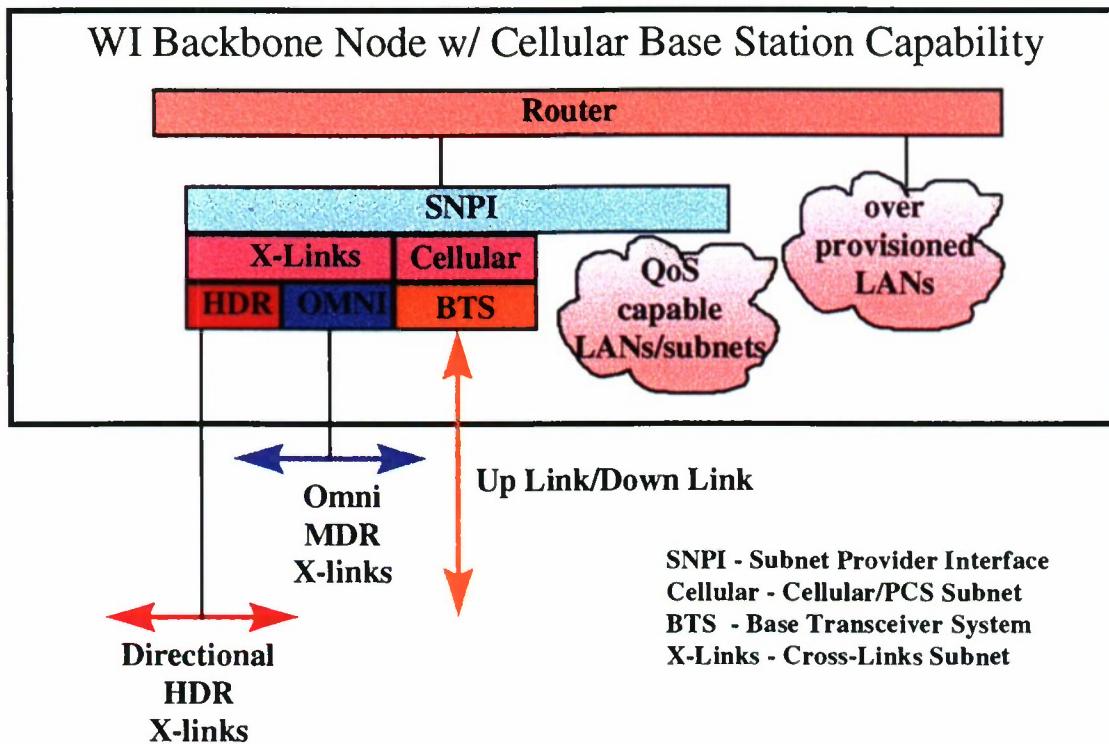


Figure 3-27. WI airborne node architecture (option #2).

Both architecture options use the omnidirectional system to point the directional antennas (Connection Management) and both support communications over the medium-speed, broadcast links. Each of the omni transmitters in the WI backbone is pre-assigned a unique frequency-hopping (FH) code. In order to receive the omnidirectional transmissions from several other Cross-Links Subnet nodes simultaneously, each node has a bank of receivers. The receivers may share a common, omnidirectional antenna. The transmit and receive antennas are positioned to minimize self-interference, and the FH patterns are designed to reduce self-interference, so as to permit simultaneous transmit and receive operation. In the ideal case, if there are N nodes in the Cross-Links Subnet, each node would have $N-1$ receivers available for listening to the omnidirectional transmissions. If there are less than $N-1$ receivers at a node, one of these receivers

should be used for scanning in order to ascertain the best set of nodes to listen to with the remaining receivers.

Some of the WI backbone nodes, such as the one on the command center's platform and those on UAVs, also contain several directional antennas and associated transmit and receive strings. These are used to support high data rate links between certain Cross-Links Subnet nodes. There are important differences in how the two candidate architectures manage communication over the high data rate links. Option 1 relies on the IP layer to manage these communications, whereas Option 2 uses the subnet layer. In Option 1, the IP routing table (and, possibly, the IP multicast routing trees) are modified by the Connection Manager when the high data rate links are changed. On the other hand, in Option 2 the Cross-Links Subnet maintains its own routing tables for communication over the high data rate links. Consequently, in Option 2, WI connectivity changes that occur using high data rate links are hidden from the IP layer. Although selecting and setting up the high data rate links is normally done automatically by the Connection Manager in both architectures, it may also be done manually if desired.

Another significant difference between Option 1 and Option 2 is the way QoS is handled. Option 1 relies on IPv6 capabilities to meet QoS needs. Option 2 uses a cell-based scheme that supports both packet switching and virtual-circuit switching. Circuit switching can be used to support voice and video QoS requirements. The Cross-Links Subnet cells are slightly larger than ATM cells. Current plans are for the Cross-Links Subnet to use 54-byte cells, which is one more byte than an ATM cell. Each Cross-Links Subnet cell has a small header that indicates the cell type. Presently, we anticipate that the following four cell types will be supported: 1) Packet Switched, 2) Virtual-Circuit Switched, 3) ATM cell carrying, and 4) Cellular/PCS cell carrying.

Since each of the subnets of Option 2 handle all routing among member nodes, the upper protocol layers (e.g., IP or ATM) are relieved of this burden. Of course, nodes external to WI that connect to it via mobile gateways still are burdened with the task of handling these dynamic connectivities. For example, if a Sincgars node connects to WI via an ACN gateway and that connection is lost and a connection to another ACN forms, then the IP layer must handle the connectivity changes. Likewise, if instead of a Sincgars node, we have an ATM node, then the ATM protocols will need to handle the dynamics of these connectivity changes. One approach to ameliorating these mobility effects is to arrange to have several nodes of the Cross-Links Subnet located on the ground (or sea). Those Sincgars and ATM nodes that were always connected to the same Cross-Links ground node could then rely on the Cross-Links Subnet to hide mobility effects.

An important feature of Option 2 is its ability to support virtual circuits among nodes of the Cross-Links Subnet, despite connectivity changes, using broadcast virtual circuits within this subnet. To do this, Option 2 routes broadcast virtual-circuit traffic over a dynamically changing, Cross-Links Backbone. Broadcast virtual-circuit cells are replicated at forks in the Cross-Links Backbone. (The Cross-Links Backbone consists of a subset of the nodes and links of the WI backbone.) The Cross-Links Backbone is reformed periodically (e.g., every 10 s or so) using a distributed algorithm. This backbone may have loops. Hence it is necessary for the Cross-Links Subnet to also provide for the proper ordering of broadcast virtual-circuit switched cells that

travel over different paths. This is accomplished by adding a cell sequence number field to each virtual-circuit switched, ATM-carrying, and Cellular/PCS-carrying cell. This sequence number is also useful for detecting cell loss.

One of the primary functions of the Cross-Links Subnet is to provide communication among the mobile base stations of the Cellular/PCS Subnet. Base-station mobility is perhaps the most significant difference between the Cellular/PCS system proposed for Warfighter's Internet and commercial cellular systems.

3.5.2.3.2 WI Cellular/PCS Subnet

The nodes and links of the Cellular/PCS Subnet are also shown in Figure 3-26. There are two classes of nodes in this subnet: Mobile Communication Devices (MCD) (i.e., subscribers) or Mobile Base Stations (MBS). These correspond to users and base stations, respectively, in a commercial cellular system. Mobile Base Stations are always collocated with a Cross-Links Subnet node. Option #2 is similar to Option #1 in that the Mobile Base Station uses a high data rate, broadcast downlink as a forward channel to the Mobile Subscribers and multiplexed, low data rate uplinks as the reverse channel. The differences in the options are in the protocols used to manage the communication.

The Cellular/PCS Subnet is a fully mobile, wireless subnet that handles its own routing. As such, it is able to hide the mobility of all Mobile Subscriber nodes from the IP Layer. To the IP Layer, the Cellular/PCS Subnet appears to be a static IP subnet which has the added benefit of being QoS capable. IP addresses can be statically assigned to each IP interface connected to a Cellular/PCS node of either type - MCD or MBS. The mapping between IP and cellular MAC addresses can be dynamically obtained via ARP (Address Resolution Protocol) or, conceivably, could be downloaded as part of a node's Complan.

MBS nodes make routing decisions based on information contained in two database structures: 1) the Global Location Register (GLR) and 2) the Local Location Register (LLR) - see Figure 3-28. The GLR maintains a current global mapping of MCD/MBS affiliations. The LLR maintains a current local mapping of MCD/cell affiliations. Whenever a handover occurs, the local LLR is updated to reflect the change. Additionally, if the handover involves a change in MCD/MBS affiliation, all GLR databases must be updated. This is done by sending a broadcast message over the Cross-Links Subnet.

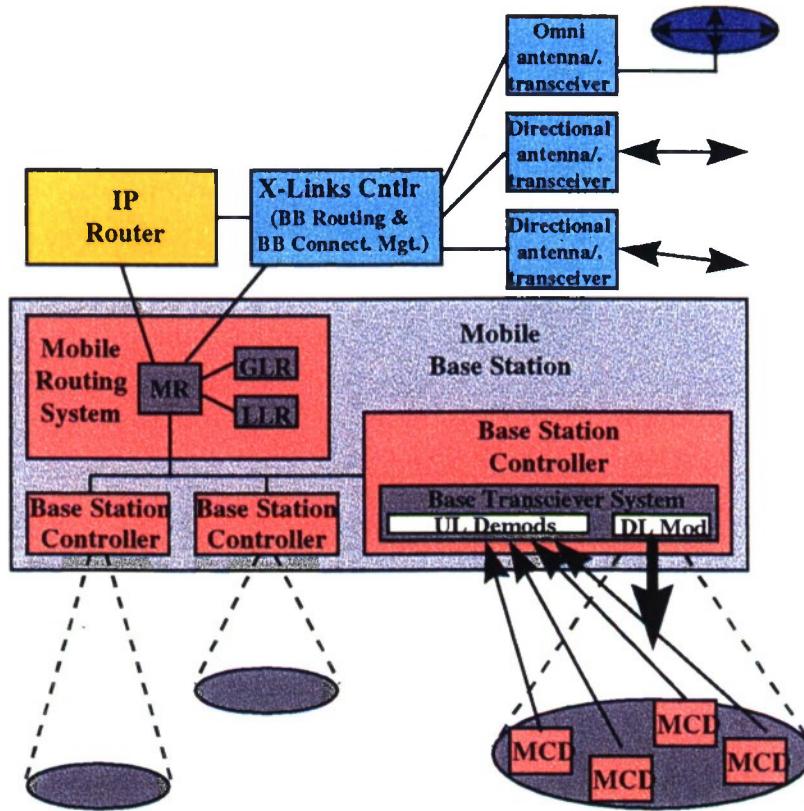


Figure 3-28. Functional diagram of an airborne Mobile Base Station (option #2) supporting Cellular/PCS subscriber links to ground-based Mobile Communication Devices (MCDs).

Both the Cross-Links and the Cellular/PCS Subnets use the Integrated Services Cell Multiplexing (ISCM) protocol to create QoS capability at the Subnet Layer. ISCM defines a Cell Multiplexing Layer and an Adaptation Layer as the mechanisms used to support integrated services (IS). In the Cellular/PCS Subnet, multiplexed flows of cells are transmitted on both the uplinks and the downlinks. The cells are 53 bytes in length so that they can tunnel through the Cross-links Subnet in the same manner as ATM cells. However, the cell headers are different than ATM cell headers and are classified as datagram (DG) cells or as virtual circuit (VC) cells. Virtual circuit cells are further classified as data or as signalling cells. VC signalling cells create virtual channels on both uplinks and downlinks for setting up and tearing down virtual circuits on those same links. VC data cells are further differentiated by a VC number. The net effect is that virtual channels are created on both uplinks and downlinks for transporting datagram traffic (one channel per link), virtual circuit traffic (multiple channels per link), and VC signalling (one channel per link). This concept is illustrated in Figure 3-29.

In many respects, this is similar to the ATM approach - i.e., using cells to create virtual channels which, in turn, support QoS commitments. One difference between ISCM and ATM is in the way the Adaptation Layer messages are packed into cells. In the ATM approach an Adaptation Layer message is always packed into an integral number of ATM cells. In most cases

this requires padding to fill out the cell that contains the message trailer. For a long Adaptation Layer message, the Adaptation Layer overhead ($[\text{trailer} + \text{padding}] / \text{length}$) will be small. However, when Adaptation Layer messages are short, Adaptation Layer overhead can approach 50%. By contrast, in ISCM, Adaptation Layer message packing is completely independent of cell boundaries. A message header (ISCM uses message headers rather than message trailers) can begin anywhere in a cell payload. In fact, the header itself can be split across a cell boundary. Partially filled cells are padded and flushed automatically as excess bandwidth becomes available; or an SNPI user may call a flush command explicitly. This approach helps reduce overhead and conserve bandwidth.

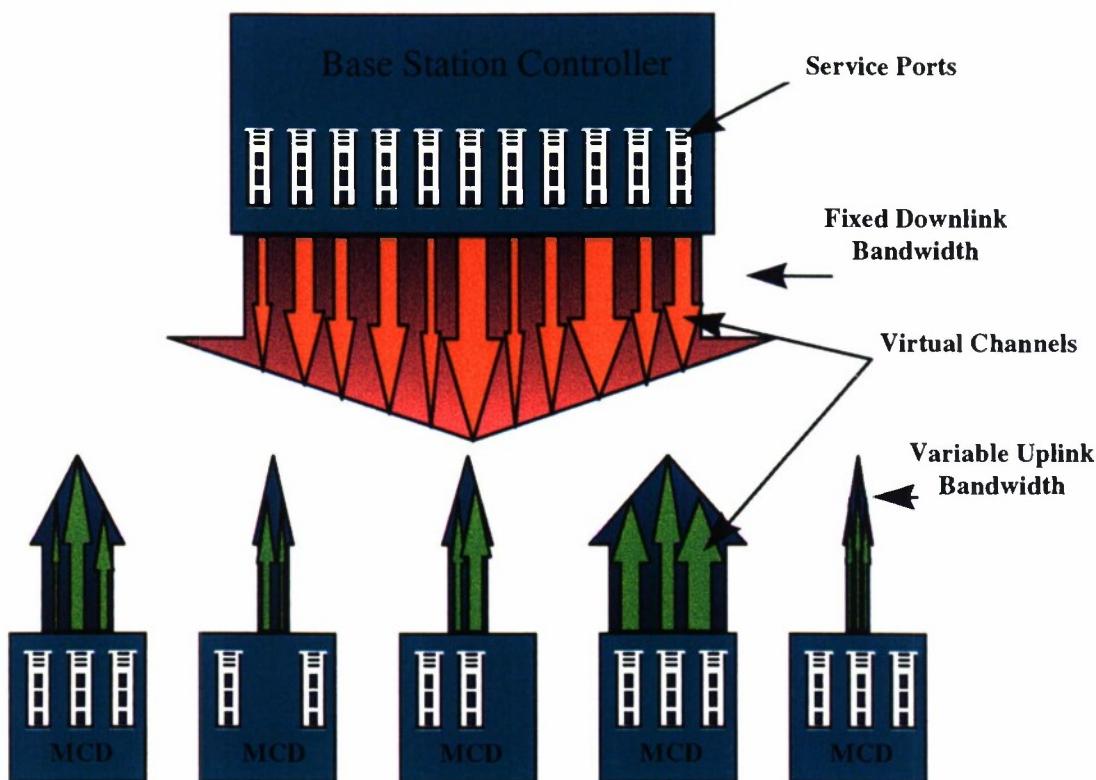


Figure 3-29. Uplink and downlink virtual channels created by cell multiplexing.

The hardware description of the Mobile Communication Device (MCD) given for Option #1 is also appropriate for Option #2. Concerns about physical size of the device, power management issues, antenna issues, etc. are the same. The differences in the options relate to the protocols used to manage the MCD. Figure 3-30 shows the MCD protocol stack for Option #2. Both options support built-in MCD applications - in particular, an IP-based voice application, a paging application, and perhaps a few simple data applications. However, at the IP Layer, the Option #2 MCD shows an IP Router. The router could be a simple, static router, or it could support a routing protocol - RIP or OSPF. Static routing can still facilitate the usage of PC

ethernet and modem cards to provide WI connectivity for a notebook computer, legacy systems (e.g., SINCgars INC, which uses PPP links), or other IP compliant systems.

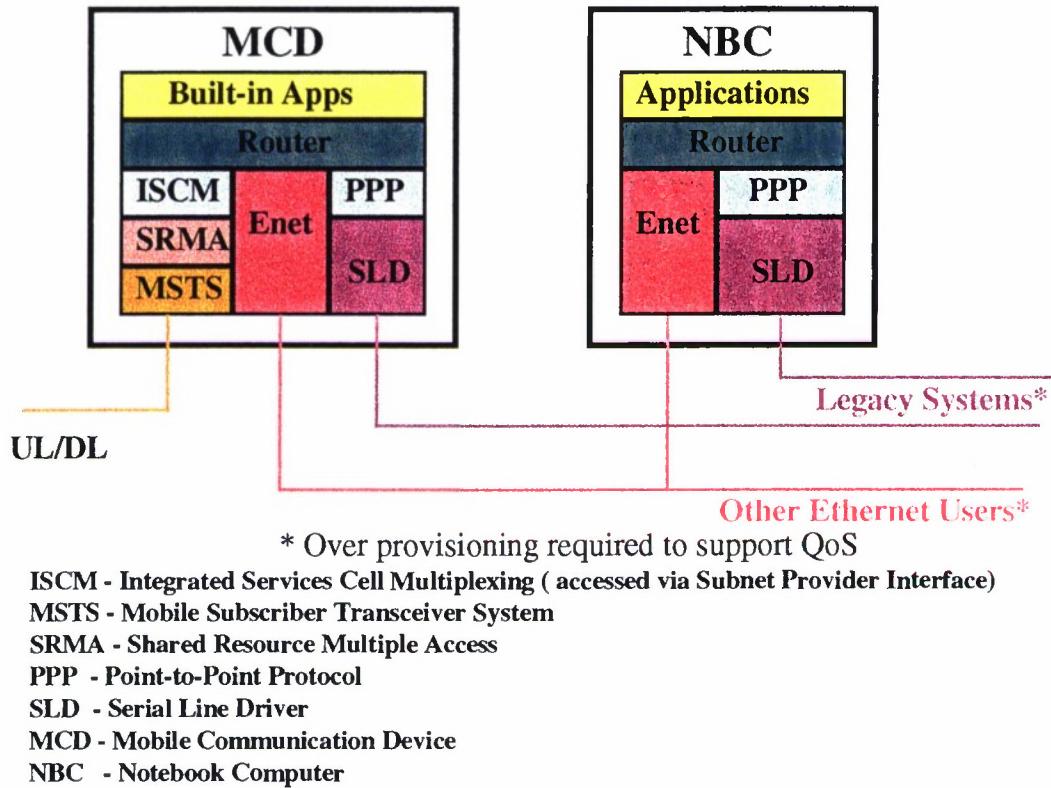


Figure 3-30. MCD and NBC configured as static IP routers for flexibility in connecting other users and subnets to WI.

The router uses the SNPI to access the WI via the Cellular/PCS uplinks and downlinks. The SNPI is layered on top of ISCM (not shown) and SRMA (Shared Resources Multiple Access), which is a demand assigned multiple access (DAMA) type of protocol for the Cellular/PCS Subnet's MAC (Medium Access Control) Layer. The SRMA protocol is actually managed by a Master Transmission Scheduler (MTS) that resides on the airborne node. Although SRMA involves both Cellular/PCS Subnet node types - MCD and MBS - the SRMA discussion is included here since it manages the MCD's MAC Layer.

Shared Resources Multiple Access (SRMA) is a methodology for dynamically allocating uplink RF receiver media resource access among a large number of multiple users based on actual traffic loading. SRMA provides the mechanisms to fully maximize access to the uplink RF receiver resources available at a base station using in-band signalling. A Master Transmission Scheduler (MTS) in the airborne Base Station Controllers (BSC) uses a host of independent decision making modules to ensure that maximum utilization of uplink RF resources is realized while at the same time minimizing the overhead imposed on the downlink data stream.

SRMA provides the mechanisms and protocols to effect the control of transmission and reception of data between the WI airborne and mobile subscribers. In the WI implementation, SRMA is combined with NRL's Integrated Services Cell Multiplexing (ISCM) which provides a single integrated data stream for a complete integrated services solution. ISCM provides offered traffic loading information, Quality of Service (QoS) commitments, integration of voice and data into a single stream of cells (data stream), reservation of network capacities to support Virtual Circuits (VC) via RSVP, and the delivery of standard UDP datagrams.

In the SRMA concept, each MCD is treated as the actual user of the uplink resources regardless of whether the MCD is serving single or multiple users. Uplink bandwidth is divided into variable size timeframes and the scheduling of uplink access is accomplished on a frame by frame basis using the actual data stream traffic available at that particular moment in time. This allows unused but reserved throughput capacity (e.g., silence during a voice communication) to be redistributed to other active (at that timeframe) MCDs.

Description of SRMA

The downlink traffic broadcast by the BSC is also broken into variable length frames each consisting of a unique Beginning of Frame (BOF) identifier that contains a unique BSC identifier, a list of Mobile Communications Device (MCD) uplink schedules, a unique Beginning of Data (BOD) identifier and broadcast data. The results of the MTS scheduling algorithms are broadcast down to all MCDs and used to regulate the time interval within the uplink frames that their transmissions are allowed to occur. Included in the list of schedules are Bid Solicitation Messages (BSM) that provide an uplink transmission opportunity to all inactive MCDs to announce their request for inclusion in the access scheduling by sending their first Traffic Transmission Requirement (TTR). The actual scheduling of BSM may occur anywhere in the current frame and may occur more than once per frame. This provides flexibility for both security reasons and for decreasing response times. Bids that have been accepted from the previous frame are announced via a Bid Acknowledgment Message (BAM) preceding the first transmission schedule for that MCD. The SRMA protocol provides an addressing mode for all commands to the MCD. The addressing modes supported are individual MCD, a group of MCDs (also used for multicast), and all MCDs. The BSC identifier within the BOF is used by the MCD to determine which BSC it is currently in communication with. The MCD can thus detect when a change occurs in BSC coverage and initiate a handoff to a new BSC if it was active at the time.

The uplink bandwidth is divided into variable length frames by the MTS. MCDs when active are allocated a slot of time within each frame to transmit based on their TTR from the previous frame or initial bid request. At each subsequent allotted time, the MCD will transmit its then current TTR along with the data that was scheduled during the last frame. If no data was scheduled for the current frame, then only the TTR will be sent. A Repeat Schedule Mode (RSM) option allows the MCD to continue transmitting on the same schedule until a new schedule is received. This feature is used to minimize downlink overhead when communications quality is high. When no data is available to be sent in the next frame, the MTS has the option to reallocate that time slot to another MCD that does have data available. This allows reserved but unused

capacity to be distributed to other active MCDs on a frame by frame basis thereby increasing overall uplink throughput while still maintaining individual QoS.

Bid Controller (BC) modules in both the BSC and MCD control all access to system resources. When an inactive MCD detects new user input data available, it will monitor the BSC's downlink frames for a BSM. At the next uplink bid opportunity a Bid Request Message (BRM) will be sent to the BSC. The BSC will perform all necessary functions required to ensure that the MCD is authorized to access the system. In addition to simple unit authorization, availability and capacity of resources requested will be verified. When all are successful, the MCD is assigned a temporary local user identification number for SRMA control. User affiliation with the aircraft and registration with the Mobile Database, etc., is accomplished as a higher level function. When a change of BSC is detected while an MCD is active, a Service Transfer Message (STM) will be transmitted to the new BSC at the time a BSM is authorized. The STM will instruct the current BSC to effect a transfer of all services from the previous BSC to the new BSC. When all users have signed off or there has been no data available from an active user for a period of time, designated in the Complan, the MCD will terminate its current session becoming inactive once again.

It should be noted that SRMA can function equally effectively with data streams other than ISCM, such as ATM, provided that they supply the traffic loading information (either on a known (in queue) or predictive basis) and the required MTS scheduling modules and Bid Controller API interfacing. Additionally, a time slice of the uplink frames could be reserved by SRMA for use by other protocols, such as ALOHA, should such a need ever arise. Additional functionality would have to be added to the BSC and MCD to support the additional protocol. However, it should be noted that using this approach may place restrictions on SRMA that might impede its ability to schedule most effectively.

SRMA Features List

The following list outlines the features of the SRMA protocol:

- Compatible with proven cell multiplexing scheme (ISCM) for managing integrated services.
- Ability to schedule on a two-dimensional basis (time/code).
- Dynamic frame-by-frame uplink scheduling allowing quick and efficient management of changing traffic loading and user response times.
- Provisions for multiple transmissions per frame.
- Flexible bid request scheduling with provisions for variable position in frame, variable frequency/code used, and multiple opportunities per frame.
- Contention for uplink access occurs only during announced bid opportunities minimizing collisions.
- Standard COTS bid contention protocols supported.
- Scheduling performed in centralized airborne nodes providing single point for upgrades.
- Scheduler designed to be modular allowing additional capabilities to be added gradually.
- Inherent automatic ranging control.
- Low scheduling requirements mean low probability of scheduling errors.

- Effective error handling.
- In-band signalling reduces hardware assets and power requirements.
- Designed for maximum possible efficiency allowing more capable hardware and software to be added as they become available without need to modify protocol or existing software.
- Protocol supports built-in beeper function.

3.5.2.3.3 IP Router Interface Segment

The IP Router provides a standard Internet interface to the Warfighter's Internet. The major effort required here is to integrate the IP Router to the WI via the Subnet Provider Interface. The Option #2 architecture can work with either IPv4 or IPv6.

Part of the integration effort involves interfacing RSVP to the Admission Control and Packet Scheduler functions provided by the Cross-Links and Cellular/PCS subnets. These functions are accessed via the QoS aware Subnet Provider Interfaces (SNPIs) supported by these subnets.

3.5.2.3.4 ATM Switch Interface Segment

There are two types of ATM network traffic that may need to traverse the WI network. The first type is ATM traffic that is destined for an ATM node outside the WI network but needs to use WI connectivity to get to the final destination. In this scenario, ATM traffic will be tunneled through the WI network to the exit point. ATM cells will get a small Subnet Layer wrapper for traversing the WI network. The wrapper is attached at the entry interface box and then is peeled off at the exit interface box. The second type of ATM traffic is the IP over ATM kind destined for a WI end user. In this scenario, the IP packets are extracted from the ATM cells at the WI network entry point interface box and then traverse the WI network to the WI end node. Both of these modes would be used in interfacing with a RAP (see Figure 3-21).

ATM switches can interface directly to the WI via the Subnet Provider Interface (SNPI) or indirectly via the IP Router (refer back to Figure 3-25). The choice of which interface to use depends on the type of service requested and the nature of the ATM traffic. If the traffic is IP (best effort delivery) over ATM, the IP packets can be extracted from the ATM cells and forwarded to WI via the IP Router. On the other hand, if the traffic is a native ATM application (i.e., it is not IP traffic) or if a stream type of service is requested (e.g., to support real-time voice or video), then the interface of choice is the SNPI. In order to implement the latter, it is necessary for the ATM switch to perform SNPI signalling to set up WI virtual circuits that are then used to tunnel the ATM cells. To properly set up the WI virtual circuits, the ATM switch must map its ATM service class (likely CBR or VBR for the case in question) and the associated QoS to the equivalent WI service type and QoS parameters. Also, it must be able to specify the WI destination address (likely the Cross-Links Subnet address) that will be the tunneling endpoint within the WI. The process is analogous to what is done by RSVP when it tries to establish a virtual circuit through one of the WI subnets. This approach is also analogous to connecting two private ATM networks through a public ATM network where the public ATM network has a proprietary Network-to-Network Interface.

3.5.3 Entry Node

There is a requirement for one active entry node as shown previously in Figure 3-18 that is the interface to the global military network. (Although this is sometimes referred to as a ground site, it could equally well be installed on a ship or be out of theater, reached over a satellite link from one of the airborne nodes.) This entry node has an active backbone link to one of the airborne platforms. Presumably, there is an alternative entry node that has a backbone link to an alternative airborne node, and which will become the active entry node in the case the link to the primary entry node degrades. It is also possible for all the airborne platforms to move out of range of the active entry node. In this case, we would need to activate a backup entry node, possibly over a satellite link.

The entry node provides many networking and interworking support functions. In Figure 3-18 we have identified a Domain Name Server (which is part of a larger hierarchical DNS service for the military). It associates an alphanumeric IP address with a 128-bit IP address. The Authentication Server is used first to verify the identity of a user requesting service and, second, depending on the user's profile, to allow the assignment of communications resources. A user's communications resource access rights are communicated from the Authenticator to the airborne basestation which then handles the real time assignments. Although not shown, there is the equivalent of what GSM calls the Equipment Identity Register (EIR) database which lists those MCDs that are not eligible for service (i.e., lost or compromised in the military context, credit denied in the commercial context). The Authentication process must poll the EIR-equivalent before allowing communications to proceed.

The ground support node also provides a special router which has many capabilities. It is a border router separating the mobile WI subnet from external IP subnets. It runs the foreign agent protocol which supports mobility, and is a central agent in the multicast routing function. It could play a central part in the interworking between the WI subnet and external IPv4 or IPv6 subnets. The details of the latter will be discussed in following sections.

Many of the entry node databases will be run on workstations connected by a LAN running a protocol such as ethernet. This LAN can also be used to support a number of local tactical database servers as well, and could mirror some external server databases as desired to reduce latency of queries from within the WI. The entry node is expected to be the major gateway point into the mobile WI. Hence the entry node router has logical connections to a network of external routers serving the tactical and strategic world.

The interworking point of voice between the mobile WI and the remainder of the tactical and strategic world is also located at the entry node. The Army uses a mix of TRI-TAC and MSE switches and the Air Force TRI-TAC switches. There is movement to upgrade both of these switch families. The WI entry node would have the responsibility of converting from the packetized voice scheme to a more conventional circuit-oriented voice session. Thus, there is a need for an interworking function. Since N-ISDN will probably be the basis for military switches in the future, then it may be reasonable to perform a packet voice to N-ISDN interworking and

allow the services to provide their own N-ISDN to TRI-TAC interworking, or any other combination to N-ISDN to military-specific switch fabric (including future ATM switches).

3.6 ADDRESSING AND ROUTING

One cannot underestimate the importance of addressing in all network design. Everyone is aware that the network address should uniquely identify an endpoint user device and hopefully an association with a person or organization. The association is an affiliated directory function. However, equally important, a destination address (in its numeric form) contains information that helps route information from the source (with its numeric address information) to the destination. For example, one's telephone number has an implicit country code, an area code, a three digit number representing one's central office and a four digit code representing a physical line from the central office to your home. Source and destination addresses are sufficient information for the network to determine optimum end-to-end routing and reserve the communications resources needed to complete the information exchanges with the desired Quality of Service.

3.6.1 IP Address Structures

IP addresses are even more powerful than telephone addresses. In this strawman architecture we are building on IPv6 technology which is the next generation IP (current is IPv4). IPv6 addresses are set as being 128 bits in comparison to the 32 bits of the earlier standard. Many of the capabilities of IPv6 networking is provided by the rich set of address types shown in Figure 3-31.

The size of the address space (128 bits) is offset by the richness of the networking structuring possibilities. In particular, we will probably exploit multicast, subnetwork router anycast, and site-local addressing for specific mobile WI implementations. As noted earlier, we can define the mobile WI as being a subnet in the larger tactical network (an autonomous system). We have included the addressing structure to indicate that some of the addresses can be used to limit certain information transfers to this subnet, partially for security reasons. The site-local and subnet-router anycast addresses are examples of these tailored addresses. We will supply more details as the multicast and security aspects of the strawman architecture design evolves. Finally, the transition from IPv4 to IPv6 will be gradual and some networks will convert before others. Consequently there must be a compatibility mode. The IPv4 computability addressing would allow interworking between the mobile WI and other tactical and strategic subnets supporting IPv4.

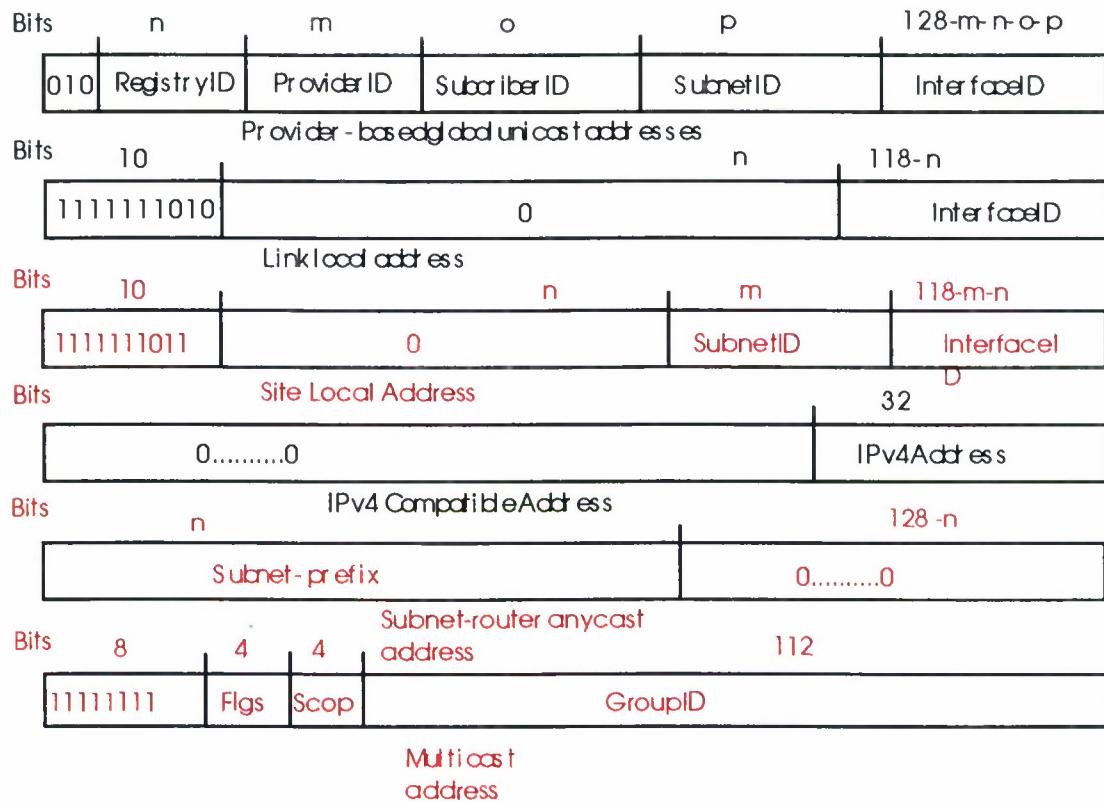


Figure 3-31. IPv6 address types.

3.6.2 IP Addresses and Mobility

As shown, IP addresses contain an identifier of the network that the user is attached to and a specific user identifier with respect to this network (here called interface ID rather than user ID since in IPv6 a user can have more than one address). For our purposes, we will assume that the user has a single unicast address of 128 binary digits. (The user can also have multiple multicast addresses but this is covered in the section on multicasting.) The desire is for the user to have a single address and any information sent to that address will reach that user irrespective of to which network the user is actively connected. This is the essence of mobility. However, the IP address has a fixed network address and this is the dilemma since the routing is determined by the network portion of the address. A solution to this problem has been implemented in cellular telephony by introducing the idea of home network and visited networks (and complemented by Home Location Registers - HLRs and Visitor Location Registers - VLRs). The analogous concepts for IP are a home agent (HA) and a foreign agent (FA), the latter the equivalent of a VLR. Since the idea of home and foreign agents are so central to routing, multicasting, and security considerations, it is useful to start with a simple example. The user is assumed to have a home network, and his IP address has this home network ID embedded in this address.

The IPv6 working group is looking at a new addressing architecture (8+8) that uses addresses consisting of a 64-bit globally unique host identifier and a 64-bit "location/routing" identifier. Unlike a conventional 128-bit address which is used to both specify a host's identity and its location in the routing topology, the 8+8 approach splits the address into two components: a static portion (~8 bytes) that uniquely identifies a host, and a possibly changing portion (~8 bytes) that specify the network to which the host is currently attached. This approach simplifies a number of problems associated with mobility and mobile security and may be applicable to the WI mobility problem.

In Figure 3-32, the correspondent host (CH) is fixed user. The foreign agent (FA) and home agent (HA) are specially augmented routers. In this example, the HA is in the MH's home network and the FA is in the foreign network where the mobile host is currently attached. When a user attaches to a foreign network, a registration process alerts the home agent of the mobile's current location.

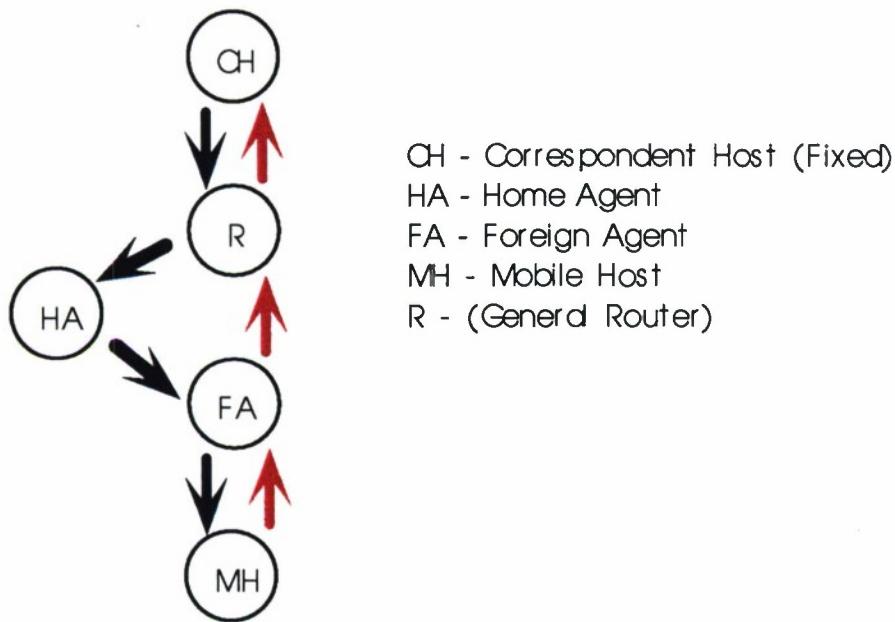


Figure 3-32. Datagram flows to accommodate mobile hosts.

Mobile IP allows mobile nodes to send "binding updates" to correspondent hosts to "redirect" future packets destined to the mobile host to be sent directly to the mobile host's new location instead of being sent to the home agent which then forwards the packets to the mobile node. In a bandwidth-constrained tactical environment, this helps reduce the traffic load on the home agent's possibly low bandwidth link and help reduce the processing load on the home agent, thereby improving its scalability. Having multiple HA/FAs provides load balancing and improves survivability. If multiple HA/FAs are used, a special mobility agent anycast address can be setup for mobile nodes to contact the mobility agent. The anycast address provides an address to

identify a service which may be handled by any number of servers for that service. The anycast address is a new feature of IPv6.

A datagram routed from a mobile host to a fixed (non-mobile) CH is not special. The CH's IP address is sufficient information for the intermediate routers to pass the datagram from MH to CH. The FA acts as a standard router.

However, a datagram from the CH to the mobile host is handled considerably differently. The CH sends the packet to the MH's home IP address. However, the MH has left its home network and is attached to a foreign network with an FA router. Consequently, the HA captures the packet, encapsulates it with the FA IP address, and sends it out. The FA captures the packet, removes the encapsulation, and routes it within the foreign network to the desired mobile host. It is seen that when the MH is a destination point, there is a two-part address translation process used to provide correct end-to-end routing. It should be noted at this time that the routing via an HA router is the root of the problem encountered when considering multicasting in a mobile environment.

There are a number of different proposals on how to perform mobile routing for both IPv4 and IPv6 for use in a global Internet. Ultimately there will be a consensus and vendors will start the code modifications needed for the Internet's routers. The choice of which proposal will be implemented depends on issues of route optimization and security. However, for the WI architecture, the actual choice is not a major issue, since the mobile WI can be made an **Autonomous System** (AS, also called a domain). The autonomous system is an accepted implementation choice for the Internet that was conceived to provide a technique for hierarchical routing and thus improve scalability for the worldwide network. An autonomous system is under the administrative control of a single organization and this is certainly a plus. Most importantly, within an AS, the choice of the routing protocol is made by the organization.

Now the key point is the following. Users outside the mobile WI who need to transfer information to a mobile user connected to the mobile WI will use whichever mobile routing standard that has been implemented to transmit information to the border Foreign Agent servicing the mobile WI. In our case the border FA is located at the ground entry node.

When one thinks of networks in the Internet sense, the common mental model is a set of organizational LANs connected to routers that are interconnected via fixed channels. Routing is performed by passing traffic from a user on a LAN to a network of routers which operate at the network address level and, according to some routing algorithm, deliver the packet traffic to the router connected to the destination LAN. Routing within this endpoint LAN merely consists of dumping the packet onto this LAN; there is no further explicit routing needed. The mobile WI is different. Once an incoming packet is presented to a border router (the Foreign Agent, FA), this agent must use the routing scheme of the mobile WI. Within the mobile WI, the routing can be a completely proprietary system that best reflects the peculiarities of the WI network.

Note also that if two mobile hosts wish to communicate with each other and they are both on the same network served by a common FA, this FA can then provide the complete routing

without recourse to their respective home agents. In the mobile WI, some of the FA functions would be distributed within the WI so that it would not be necessary to include the main FA function on all the paths connecting users in the WI. A more detailed overview of a possible (mobile WI) intradomain routing strategy will be given in Appendix E.

The registration process of a mobile user on the mobile WI is a key step in the mobility scheme. Not only does it establish the location of the connection point relative to the airborne platform, but it is used to start the foreign agent to home agent exchanges to update mobility databases. Periodic registration is a fundamental process in any mobile cellular system. We will detail more of the registration process for the mobile WI in the various appendices that address mobility and multicasting.

3.6.3 IP Addresses and Multicasting

It is useful to review the reason why multicast is useful and the basic mechanisms supporting multicast. This is first done without considering mobility which greatly complicates the situation. Note that multicast is an integral part of the IPv6 architecture.

Multicasting is meant for UDP packets which is best effort delivery. It can be used very effectively on links with low bit error rates. For lossy links, use of UDP should be considered on a case by case basis to determine whether it is really appropriate for the application intended.

All are familiar with broadcast services. Broadcast services are heavily used in TCP/IP networks. For certain tasks, broadcast uses more communication resources than necessary. There may be many net members who have no interest in the information being transmitted yet valuable capacity may be used in getting the information to every endpoint even though the endpoint does not read it. Thus the basic idea is to organize users that share a common interest into a multicast group. This group would have a single, common address. Earlier the IPv6 addressing structure was shown; the multicast IP address has no network identifier and this is appropriate since members of a multicast group may be located anywhere, even in different networks.

A single multicast addressed packet substitutes for creating multiple unicast IP addressed packets to the group members. Multicast eliminates transmissions over those links not needed for routing information from source to the set of active destination points. When a multicast member is not active, the message is not directed towards this member. Clearly, multicast is more bandwidth efficient than the general broadcast flood techniques. The trend is to use multicast capabilities as a substitute wherever feasible. However, one must recognize that implementing broadcast is considerably simpler.

There are a number of proposed multicast routing strategies proposed for fixed networks. Each proposed multicast routing algorithm builds the multicast trees in different ways. The routers must be multicast-capable (meaning that they can build appropriate multicast trees). One of the multicast routers has an additional function. It must act as the manager of a multicast group. It contains a database of active membership of a multicast group, and group members

explicitly communicate with this to indicate joins or leaves. This is an IGMP function in IPv4 and an ICMP function in IPv6. One multicast strategy (the Distance Vector Multicast Routing Protocol, DVMRP) is for each router to keep track of the best path to the source of the multicast packet. Then when a multicast packet arrives on the best path link for that particular source, it is sent out on every other connection of that router. If the packet came on an unexpected link to the router, it is dropped. This dropping is what limits the flooding implicit in broadcast.

In the mobile WI we must consider mobility and multicasting simultaneously because of their interactions. Mobile IP (IPv4) is well along in the standards process in that a completed draft standard is available, while Mobile IP for IPv6 is still evolving. The Multicast standards for IP have been mainly concerned with fixed networks in mind. The types of multicast routing protocols considered reflect this bias. Even when wireless links were considered, only point-to-point wireless links are referenced. When the user at the end of this point-to-point link leaves, the link is dropped from the multicast route calculations. The standards groups have not considered the broadcast nature of radio links where unaggregated users listen independently to a single RF transmission. In this case, the packet would be eliminated from the RF transmission only when all the active members have left the multicast net.

Multicast coupled with mobility is not yet a standard and there are many important outstanding issues. A fundamental issue is concerned with the way that multicast routers compute routing trees from each source member to the active listening members. However, mobile users are addressed by their home IP and consequently, in some way, one has to consider the role of home agents and foreign agents since they create different routes from what the multicast routing expects. This results in legitimate datagrams being rejected as they enter a multicast router, assuming one uses the current multicast routing algorithms, as well as the more obvious inefficiency of routing messages first to the Home Address only to have to reroute them to the Foreign Agent. Instead, the multicast tree should be constructed with complete knowledge of the mobile (Foreign Agent) address of every member of the multicast group.

Since multicasting and mobility together are not yet a standard, and given the proposal to make the mobile WI an autonomous system, a proprietary scheme can be used within the WI if it is made sure that, for multicast members outside the mobile WI, a standardized multicast distribution scheme is provided. Some of the possibilities are addressed further in Appendix F.

3.7 SECURITY

The goal of the mobile WI is to provide the untethered Warfighters with secure information access possibilities that they never had before. The information exchange between communicating entities (human to human, human to machine, and machine to machine) will be independent of the geographical separation. The Warfighters can be located anywhere (in friendly or unfriendly territory) and still be able to communicate in a secure manner. Moreover, capture of an MCD and the Warfighter can be anticipated, and thus we need to set up mechanisms to contain the effects of compromise.

We believe that the traditional military security architectures are unwieldy and unaffordable when applied to the untethered Warfighter. There are a number of commercially driven developments that show great promise in being able to be applied to support military operations. Since we selected a connectionless information infrastructure and have targeted IPv6 as being our fundamental facilitating networking technology, we will see how we can optimally exploit its built-in security enabling mechanisms. Our goal is to exploit these commercial ideas as much as possible and to assess how military security is evolving (or can evolve) to fit the security architecture that we are using as our overall strawman WI architecture. In particular, we will start the process of assessing how the NSA Fortezza products can be utilized. The promise of Fortezza is in its low device cost.

Since spread spectrum technology will probably be used as an integral part of our architecture, it is possible that it can provide a TRANSEC function that provides a satisfactory measure of link level security. This frees us to concentrate on higher level security mechanisms and the remainder of this section is concerned with network level and above security issues.

It is easy to associate security with encryption devices, but these devices are only a small part of the total security problem. In a comprehensive WI security architecture, one must consider the full spectrum of security services. Figure 3-33 shows a complete security taxonomy.

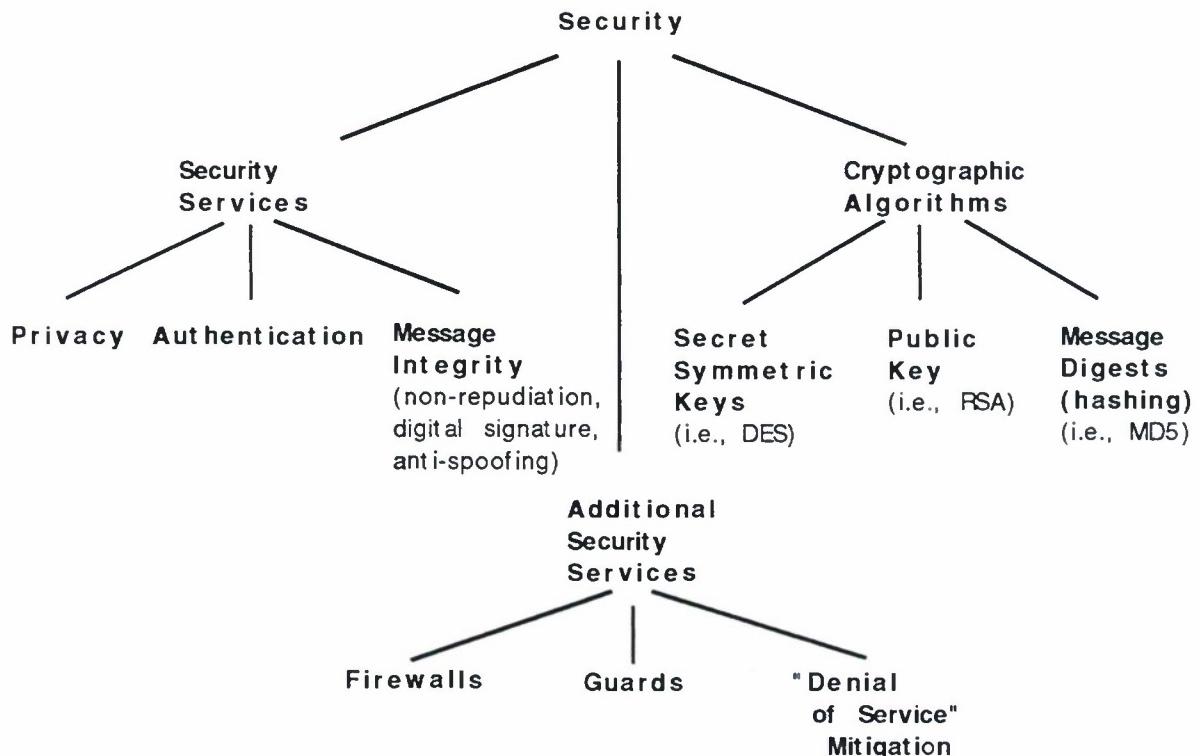


Figure 3-33. Security taxonomy.

In this taxonomy we have separated the general security services from some more unique services. **Privacy** refers to the exclusive ability of only legitimate receiver(s) to read the transmitted information. Everyone else sees ciphered text at best. **Authentication** refers to certifying that the source of the information is who he claims to be. **Message integrity** refers to the fact that the message has not been changed in transit between the sender and the receiver.

Cryptographic algorithms are the mechanisms used to provide all these general services. All three types are candidates for our security architecture, although secret symmetric keys have been the basis for all military security systems.

Firewalls are special implementations reflecting the fact that it is extremely difficult to build perfect and homogeneous security systems. They are an attempt to concentrate some of the shortfalls into a system that can be more easily controlled. **Guards** are another implementation set that prevents higher-level security messages from leaking onto lower-level security networks (but not the reverse). **Denial of service mitigation** is intended to prohibit an intelligent foe from denying a legitimate user from using the communication and networking resources. In the mobile WI we will examine only firewalls and denial of service mitigation.

At this point we do not have a WI security architecture but rather we have scoped out a direction that we intend to follow to provide a consistent security architecture. Moreover, we intend to embed the security elements in the strawman design and not add them at the end of the design. The latter approach invariably ends up with security compromises or greatly increased costs to accommodate costly but somewhat mismatched security products.

Overall, for the proposed mobile WI, we are advocating an advanced IP network realization based on IPv6 technology that is considerably advanced over the current military networks. This IP technology is evolving from the commercial world needs with respect to both mobile communications in general and global Internet security concerns. In fact, the implications of both mobility and security are what are driving the design. We will see that the IPv6 security architecture is “open” in the sense that different authentication methods, traffic encryption methods and key management systems can be accommodated. Our goal will be to assess how the IPv6 security framework and specific NSA security products can be merged into an affordable security solution.

3.7.1 Security and IP (Historic)

Secure IP networks have been a key technology in the strategic networks for two decades and are being extended into tactical networks today. SIPRNET is one such example. The security of these networks is an overlay and does not change the basic IP protocols. The supported IP protocol used is IPv4. The important point is that the version of IPv4 used has the following characteristics:

1. The IP protocol is such that routers are not involved in any security issues.
2. The IP protocol has no built-in security support features.

We do not mean to imply that security features cannot be added to IPv4 protocols. In fact, there is standards work to provide this modification. Security support will show up as option fields in the IP header. However, security additions to IPv4 are add-ons and, like most add-ons, are not an optimal strategy. However, there will undoubtedly be secure IPv4 protocols and the mobile WI will have to interwork with them to the maximum extent.

Rather than add security to the earlier IP protocols, NSA commissioned network level security device development for commercial systems starting almost 15 years ago. This program was intended to provide security at the network level, specifically on the WAN parts of extended WAN/LAN networks. NSA's Secure Data Network System (SNDS) program initiated development at a number of commercial mainframe and minicomputer companies. A parallel military network level security program resulted in a Motorola system called the Network Encryption System (NES) which incorporated many of the lessons learned in the SNDS program. One important goal of this program was to inhibit traffic analysis on the WAN backbone, and traffic analysis is a network level security issue. It should be noted that with the NES devices the standard IPv4 was not modified so the normal WAN routers would not know that it was handling secure traffic.

One problem is that the NES devices are expensive and they only hide the traffic analysis on the WAN segments. LAN segment headers were in the clear. Both of these are valid reasons for not adopting the NES network security approach. However, it still is useful to provide some sort of security at the network level.

3.7.2 Voice Security (Historic)

Secure voice will be an integral information type in the tactical world and will remain so within the mobile WI. For the tactical world, two popular secure voice systems that operate over military switches are DSVDs and STU IIIIs. These devices have encryption at the user level. There is no network level security for voice, and hence link level encryptors have been used. It should be emphasized that the mobile WI will not use either DSVDs or STU IIIIs. There is also a program called STE (Secure Terminal Equipment) to produce the next generation secure voice equipment. This program is also not part of the secure voice solution that we are proposing for the mobile WI. However, we still need to see how these systems can interwork.

In the mobile WI we intend to use packetized voice and to some extent treat voice in a similar manner to other forms of information, including security. This means that we could use the security enabling IPv6 mechanisms discussed next.

3.7.3 Key Management

Key management is a huge problem in traditional secure military networks. There is just too much keying material that needs to be managed. Typically, it requires a considerable physically secure infrastructure for key generation, storage, distribution, and injection into the end systems.

For the untethered user, key management can quickly become a problem. For the untethered user whose location is unpredictable (with respect to access to a key management center location), one desire would be that session level keys be derivable dynamically in the user device and in their correspondent end point entity. This is the operational mode of most planned commercial mobile systems. Similarly, there are certain parts of the overall security architecture where public key systems can be used.

It must be noted that IPv6 does not specify a Key Management protocol, and there is nothing in the IP security option headers (AH or ESP) that carry key management data. If it is to be carried, it will probably be carried over higher level protocols (TCP/IP). There are many secure ways to cooperatively generate shared keys. For certain applications, it will be necessary to evaluate novel key management mechanisms where dynamic session keys can be efficiently generated as needed.

3.7.4 IPv6 and Security

There appears to be no clear choice on the best technology on which to base a security architecture for the WI. However, since we believe that the Internet will be the main networking technology thrust, at least for the next 10 years, and that it is evolving towards IPv6, we feel most comfortable in selecting the security support mechanisms in IPv6 as the starting point of our security considerations.

The advantage of IPv6 is that security is embedded in the overall protocol design. While most of the IPv6 developments are conditioned towards the commercial and government world, there is little reason why they cannot be adapted to satisfy a variety of military uses. As noted, IPv6 is “open” in a security sense, since it does not mandate specific encryption algorithms either for authentication or for encryption. However, there are still standard default algorithms (keyed MD5—hashing and DES) in case no alternative is provided. Default algorithms may have some use in secure communications between WI network elements.

Since we are looking at an entirely connectionless system at the network level, it makes sense to apply the IPv6 security model to voice and video packets in addition to data packets. This is one fundamental difference between this approach and the traditional voice versus data security approaches used in the military. Our approach will be to assess how far IPv6 can directly support our mobile WI security desires and then judge where the shortfalls are. Looking at the limitations of commercial security is useful since one anticipates that this environment will not accept costly, special purpose security hardware and will look for solutions that simplify key management and security operations.

To understand IPv6 security it helps to start with the IPv6 packet structure. This is shown in Figure 3-34 (which also shows IPv4 as a comparison).

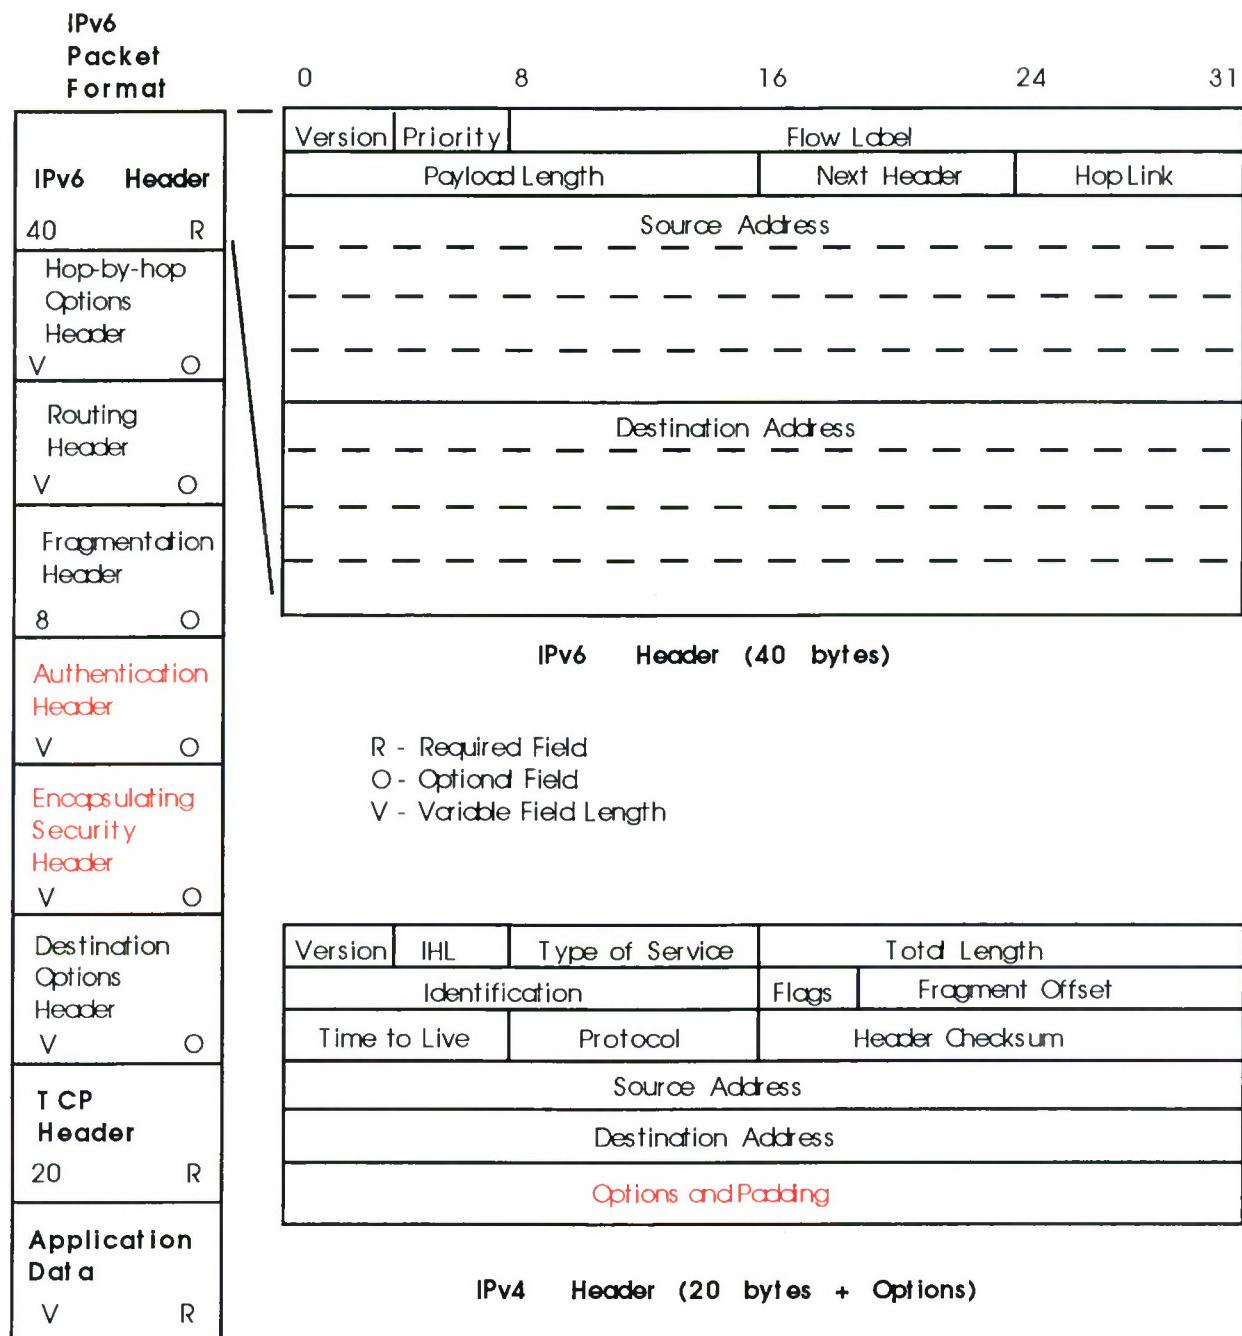


Figure 3-34. IPv6 header with optional subheaders.

This figure shows a basic comparison of the packet header structures of IPv4 and IPv6 (including subheaders). It is seen that only IPv6 has subheaders for their options while IPv4 has its options as part of the IPv4 header. The security-related parts of IPv6 are the Authentication Header and the Encapsulating Security Payload Header. Security is an integral part of the design of IPv6, whereas it is an add-on to IPv4 in which it would be in one of the option fields.

Only IPv6 will be concentrated upon because it is the future of IP technology and is being designed to better handle real-time information flows. One can express alarm at the length of an IPv6 header, but one has to trade off this length for the added capabilities that the header and its optional subheaders provide.

The first security related subheader considered is the Authentication Header shown in Figure 3-35.

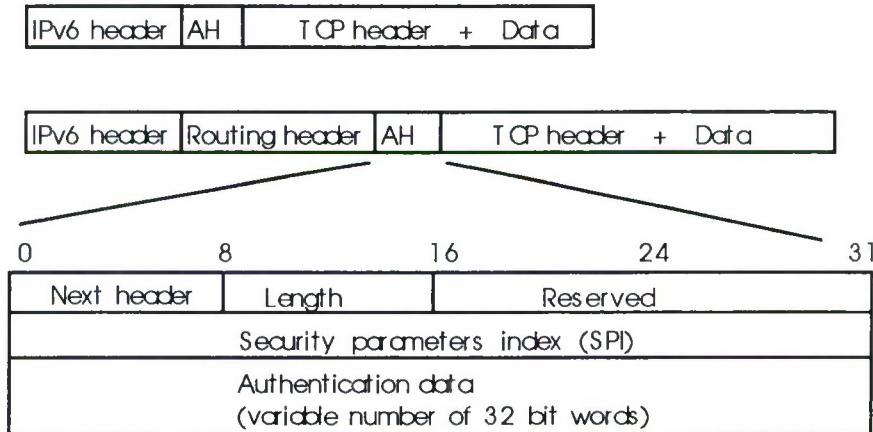


Figure 3-35. IPv6 Authentication Header (AH).

At this point it is useful to define some terms that are critical to understanding IPv6. A security association is a set of agreements between a sender and a receiver. A specific security association would agree on keys to be used for authentication and encryption, authentication algorithm, encryption algorithm, security level, key lifetimes, etc. Each security association has a 32-bit Security Parameter Index (SPI) which is a pointer to the security association fields to be used for the communicating entities.

It first should be recognized that authentication is a process of guaranteeing that the source of some data is who he claims to be. Authentication is aided by the Authentication Header (AH). Authentication does not mean that the user traffic need be encrypted. Authentication data is typically some non-invertable checksum of a combination of some data traffic plus a secret key plus some fields in the header not subject to change as the packet moves from node to node. The authentication data in this subheader need not be encrypted. An agreement on checksum by the receiver operating within the same security association prevents certain spoofing attacks against the network and validates the user as being the source of the data.

We also have a common security association for multicast groups. In this case, one member of the multicast group selects the SPI (and hence security association) for all the members of the multicast group.

There is tremendous flexibility enabled by using security associations. Different communicating entities may have different communications needs. For example, one could enable secure user-to-user traffic with a Fortezza product card, and for router-to-router communications use the default commercial security mechanisms. The choices can be made on the basis of implementation cost and difficulty, required security levels, processing rate requirements, etc.

If it is desired to encrypt the user traffic, then the Encapsulating Security Header option is used. This is shown in Figure 3-36.

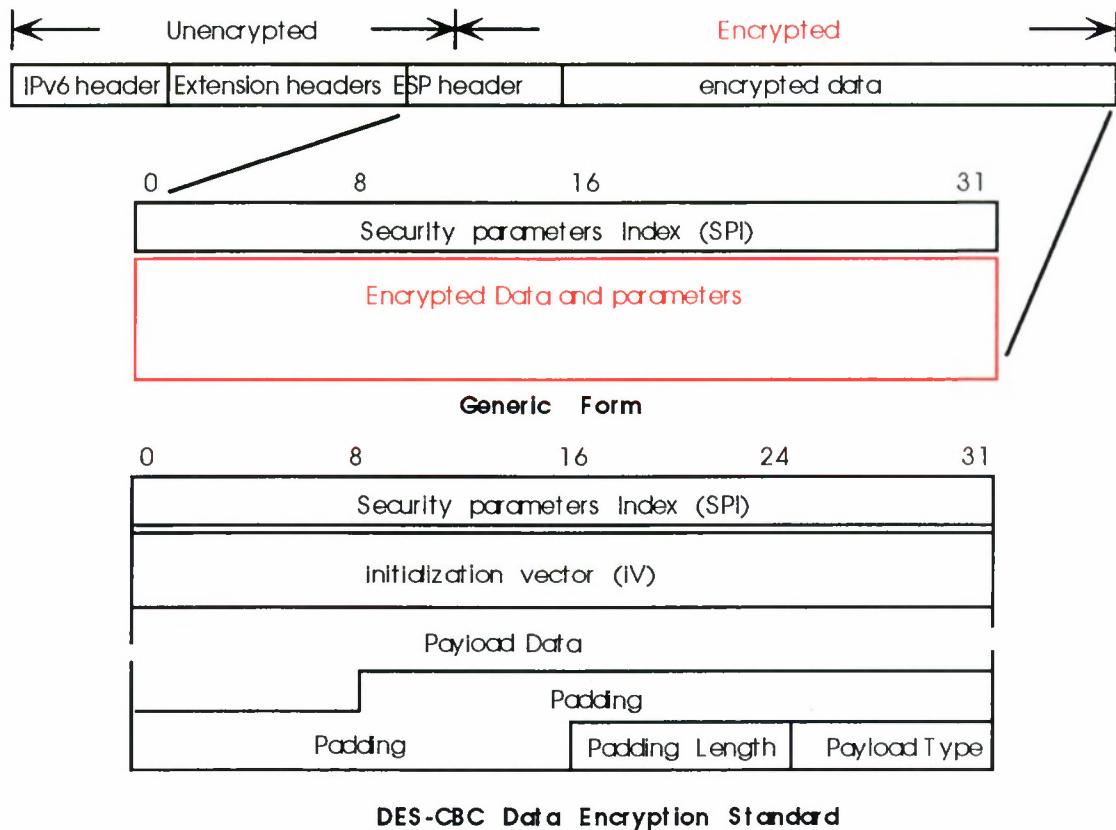


Figure 3-36. IPv6 Encapsulated Security Protocol subheader.

In most military applications it is imperative to encrypt the data. For IPv6, we make use of the additional Encapsulating Security Payload (ESP) subheader. Some header extensions can be placed before the ESP header. One should note that the ESP subheader itself is partially encrypted and everything following, including user data, is encrypted. Only the SPI is in the clear. This is necessary to mutually agree on a security association. It should be noted that depending on the type of security desired, the Authentication Header (AH) can precede the ESP and be in the clear or follow the ESP and be encrypted.

The figure also shows an example where DES is being used as the encryption algorithm. The SPI and the security association will define which authentication and encryption algorithm to use; they need not be the same. This is important, since some algorithms may need to be performed faster and could use a simpler computational algorithm.

As in the case of authentication, the ESP header can be made to work with either unicast or multicast traffic. This is clearly an important attribute.

It should be observed that we apparently have ruled out the SSL (Secure Socket Layer) approach, which is being advocated by a number of commercial companies. We feel that SSL is too proprietary and oriented to specific vendor approaches in the commercial world. We will reexamine this position after we have determined what the issues are in building an API that is responsive to the special needs of this wireless environment.

3.7.5 The Potential Role of Fortezza

As a first impression, Fortezza appears to be a PCMCIA card (now called simply PC Card) that is meant to provide inexpensive user-level encryption suitable for use in PCs and in certain workstations and servers with network attachments. In fact, the target network is probably the Internet.

Fortezza has a considerable history. It has always been a key component of DoD's Defense Messaging System (DMS) which is intended to replace AUTODIN. Unfortunately, the program names have changed frequently. In 1991, the program was called the Pre-Message Security Protocol. In 1993, the program was called MOSAIC with the device (called a token) named the Tessera Crypto card. It was intended to provide direct end-user to end-user security services to Sensitive but Unclassified (SBU) electronic mail (SMTP and X.400). MOSAIC consisted of a TESSERA card and software drivers which interfaced with a commercial E-mail package that accessed the TESSERA card. The TESSERA card conformed to a PCMCIA type 2 standard and hosted the CAPSTONE chip. It was intended for purchase at the end of 1994.

In 1994, this effort was merged into the wider MISSI program and the card was renamed the Fortezza Crypto Card.

In April of 1993, the government (through NIST) introduced a new encryption initiative aimed at providing a high level of information security but allowing, at the same time, the means for designated government organizations to have access to the information. The organizations with access can represent law enforcement, public safety, and national security. In February of 1994, the government announced adoption of this initiative as the Escrowed Encryption Standard (ESS). The ESS is a voluntary government standard for Sensitive but Unclassified information sent over phone lines, analog or digital, and encompassing voice, fax, and data.

The initiative rests on a special tamper-resistant hardware encryption device (the Clipper Chip is one example) and a Key Escrow System (KES). This combination gives the government access to a Device Unique Key that enables the decryption of all communications encrypted by

the chip. This key is generated and programmed onto the chip after the chip is manufactured but before it is placed within a security product. At the same time, the key is split into two Key Components which are encrypted and provided to separate Key Escrow Agents. Under legal authorization, a government representative needs to acquire both Key Components which are in turn combined in a special Key Escrow Decrypt Processor to obtain the Device Unique Key that can decrypt the intercepted communications. The agencies selected to decode this information will be determined by the Attorney General. Initially, the KEAs were intended to be NIST and the Department of the Treasury Automated Systems Division. NSA was the developer and the FBI is the initial law enforcement user.

There are a number of chips that will support this capability, one of which is the CAPSTONE chip. The CAPSTONE chip contains a Secure Hash Algorithm (FIPS 180-SHS), the Digital Signature Algorithm (FIPS 186 DSS), and a Skipjack type II CBC algorithm which replaces the DES and a type II key exchange algorithm (KEA). The chip also contains a high-speed exponentiation algorithm and a random number generator based on a pure noise source. This chip is intended for both government and commercial use. The Skipjack encryption algorithm transforms a 64 bit input block into a 64 bit output block using an 80 bit secret key. This is a symmetric key system. It replaces the DES but has a key which is 24 bits longer. This can be operated in one of four modes:

1. Electronic Codebook (ECB)
2. Cipher Block Chaining (CBC)
3. 64 bit Output Feedback (OFB)
4. 1, 8, 16, 32 or 64 bit Cipher Feedback (CFB)

Both the Skipjack and LEAF algorithm (to be described) are classified; the former to prevent a design copy without key escrow, and the latter to prevent products to be built that would correctly interoperate with ones that have the official correct key escrow function. LEAF (Law Enforcement Access Field) is an embedded algorithm that generates a 128-bit field that is transmitted with all encrypted communications and provides a mechanism for securely transmitting the encryption key for the session (KS) for use by law enforcement access.

However, this is only for the law enforcement official and not for the distribution of KS (session key) for the users (the receiving user cannot get the KS from the LEAF). The distribution of the session key, KS, for the users is done separately.

The desire of the government is to export this concept to the commercial world. However, for the commercial world, the government would make use of the built-in key escrow system to monitor selected traffic. Under the popular name of Clipper, this entire system has generated tremendous controversy in the US and international community. However, this does not detract from its potential as a device to support military operations.

Let us review its attributes. It will be inexpensive (approximately \$100) and available from a number of vendors. It conforms to a standard type 2 PCMCIA interface and is both lightweight and requires low power. All of these are important attributes for the Warfighter

MCD. While the initial cards are meant to support sensitive but unclassified, there will be newer chip versions that can support up to Top Secret Level.

Even more important, Fortezza is not limited to providing user traffic level encryption. It can support data integrity (anti-spoofing), originator authentication, non-repudiation and data privacy. Our immediate task is to see how this can be done in the context of the current strawman which is initially considering use of the security support features of IPv6.

At this point the capabilities of Fortezza have not been evaluated by the WI community. It is anticipated that multiple, simultaneous secure end user sessions can be maintained (up to 10) and this is desirable.

It would be useful if one could program the Fortezza chip to produce a secure key that could be imported into the spread spectrum system as an SS key. This desirable feature is incompatible with the design of the Fortezza concept which does not reveal any secure key. However, one could generate a key separately through some reasonable pseudo-random key generation software algorithm in the base station and then use this as data and ask the Fortezza card on the airborne subscriber to transmit this down to the subscriber handset's Fortezza card where it would be taken off as data and passed to the SS unit as the session key for the SS/TRANSEC unit. Note that there are separate keys for the uplinks and downlink generated and distributed in the same way.

3.8 NETWORK AND SYSTEMS MANAGEMENT

The WI network topology consists of clusters of radio networks interconnected using Airborne Relay Platforms. Each cluster needs to be in contact with other clusters and other external networks. The goal of the network management system for mobile WI is to perform monitoring and control of network devices, systems, and applications (wherever possible) in a near-optimal fashion to match the available communication assets with the critical communication needs. This must be done in a manner that will accommodate topological changes in a timely and efficient fashion and is responsive to user needs.

3.8.1 Management Protocols and Architecture

SNMP is universally accepted as the protocol for managing computer networks. In traditional management using SNMP there is one agent per network device and a centralized manager. The agents access information (in the form of variable values) about the managed devices and make it available to the network management system. Each of these variables is referred to as a managed object. All managed objects are contained in the Management Information Base (MIB), a database of the managed objects. There are two types of communication between the agents and the manager. One type of communication is initiated by the manager when a request is made to view or change the value of a device parameter by sending a get/set message to the agent. The agent executes the request (if possible) and returns the current value of the parameter (in the case of a get message) or the new changed value of the parameter (in the case of a set message) to the manager. The other type of message is initiated by

the agent when a device parameter value falls outside a pre-set range (usually pre-set from the manager). In this case the agent sends a message (known as a trap message) back to the manager indicating the condition and the new value of the parameter. There is no other intelligence in the agent software. All of the control resides within the network management software. By monitoring appropriate device parameter values and trap messages, the network manager determines what control changes are needed at the devices and executes them.

The WI network has significant differences from a regular terrestrial IP internetwork. In the terrestrial IP world the routers and switches are statically positioned most of the time. The physical connectivity is usually stable. The data loss in the IP world is also minimal because of the noiseless operational environment. None of this holds true for the above wireless network. Another major difference is the bandwidth available for transporting management traffic between the Airborne Platforms and a centralized management center. Constant monitoring of the Airborne Platforms and other managed components from a centralized management center and making decisions based on such monitoring is not an option due to the bandwidth limitations. The objective could be achieved by distributing the network management function as much as possible to the relays while maintaining overall control from a centralized network management center. Distributed multi-level network management is well suited to this network where a mid-level manager/dual-purpose agent is present at each of the Airborne Platforms interacting with management agents on network components under its domain and with the main network management station at ECOC.

Employing mid-level managers allows for distribution of management functionalities where control is delegated from one management station to another. At present, partial implementations of this type of an architecture have been implemented and deployed in the Internet using proprietary architectures by COTS NM vendors. To promote interoperability as well as a common framework upon which various systems can be built, the DISMAN working group (under NM directorate of the IETF) is currently working on a distributed management standard. An Internet standard is not expected for a couple of more years. The WI management system development team is in a good position to interact with the standards effort in the DISMAN working group so that the standards would satisfy WI network management needs.

3.8.1.1 Implementation Schemes

If the SNMP agents running on the airborne platforms could be extended in capability so that they could make the simple and straightforward control decisions (at the least do event correlation and data aggregation and pass only the concise information to the main network manager), the network management traffic between the centralized manager and relays could be minimized. In this scenario, the agents running on the relays are not passive entities; they take an active role in management. The centralized manager only makes the more difficult control decisions; the ones which need the global knowledge available at the centralized control center.

One scheme for management is as follows: at initialization, the Control Center/ECOC Network Manager downloads a specific configuration to each of the airborne platforms using SNMP or some other mechanism. The airborne platforms each run an SNMP capable dual-role

entity (manager and agent) which monitors the performance metrics for the network components for which it is directly responsible. When the local agent on an airborne platform notices that a pre-set fault-management limit condition/metric is exceeded, a message is sent to the control center indicating the same (SNMP is ideally suited for this via trap messages). The control center then decides (since it has global information available) whether to download a new configuration file to each airborne platform or to the specific airborne platform that exceeded the limit.

In another variation of the above scheme, each airborne platform is pre-loaded with all the possible configurations and the Control Center/Network Manager decides which one to follow depending on the run-time loading. This approach eliminates the need for downloading different configurations during run-time. In this scenario, an SNMP set message can be used to set the airborne platform to the new configuration when a metric is exceeded in one of the gateways.

Possible Extensions: The above scheme can be further extended to semi-automated management of the mobile networks where the final control rests with the centralized control center while each of the gateways function as semi-autonomous managers which function within well defined parameters.

3.8.2 Network Management Functions

The network management functions to be supported for the WI are Fault Management, Configuration Management, Performance Management, and Security Management.

3.8.2.1 Fault Management

Fault management involves correlating events/traps received from the network, probing for additional information from the components if needed, running diagnostic tests on the network if needed, and isolating and resolving problems.

3.8.2.2 Configuration Management

Configuration management involves monitoring and controlling the configuration of devices, systems, and applications.

3.8.2.3 Performance Management

Performance management measures traffic flow across the network, calculates measures such as the number of packets/cells that are successfully transmitted against those that are dropped, etc., and implements configuration changes to optimize performance/efficiency.

3.8.2.4 Security Management

Security management deals with controlling access to network resources through the use of authentication techniques and authorization policies. A secure network management protocol with encryption and authentication capabilities is required to achieve this function. SNMPv3 may be the first version (to attain official standard status) which satisfies this requirement.

REPORT DOCUMENTATION PAGE

*Form Approved
OMB No. 0704-0188*

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)			2. REPORT DATE 28 January 1998		3. REPORT TYPE AND DATES COVERED		
4. TITLE AND SUBTITLE Architecture and Concept of Operations for a Warfighter's Internet, Volume 1			5. FUNDING NUMBERS C — F19628-95-C-0002 PR — 602				
6. AUTHOR(S) Edited by MIT Lincoln Laboratory							
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Lincoln Laboratory, MIT 244 Wood Street Lexington, MA 02173-9108			8. PERFORMING ORGANIZATION REPORT NUMBER				
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) DARPA/ISO 3701 N. Fairfax Dr. Arlington, VA 22203-1714			10. SPONSORING/MONITORING AGENCY REPORT NUMBER ESC-TR-97-064				
11. SUPPLEMENTARY NOTES None							
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.				12b. DISTRIBUTION CODE			
13. ABSTRACT (Maximum 200 words) <p>Military operations in the twenty-first century will be conducted in an increasingly information-rich environment. But delivery of this information is difficult in the forward areas of the tactical theater where current communications equipment is slow to deploy and not matched to the mobility of forward forces. This report describes a "Warfighter's Internet" that provides responsive data and voice communications to individual warfighters using hand-held cellular-like handsets that have a wireless connection to a backbone network of airborne communications nodes that are within line of sight of theater forces, within line of sight of each other, and are also connected to command and support centers through the Global Grid. The airborne nodes self-deploy and the handsets arrive with the users, so the network can be available immediately. This report describes design approaches to meeting the technical challenges in this system, which lie in the adaptation of Internet data protocols to the mobile environment (commercial cellular systems do not have to deal with both mobile users and a mobile backbone) and to the efficient use of wireless capacity for bursty data transmission.</p>							
14. SUBJECT TERMS wireless communications data networking					15. NUMBER OF PAGES 119		
					16. PRICE CODE		
17. SECURITY CLASSIFICATION OF REPORT Unclassified		18. SECURITY CLASSIFICATION OF THIS PAGE Same as Report		19. SECURITY CLASSIFICATION OF ABSTRACT Same as Report		20. LIMITATION OF ABSTRACT Same as Report	